

Recurrence Width for Structured Dense Matrix Vector Multiplication

ALBERT GU*

ROHAN PUTTAGUNTA*

CHRISTOPHER RÉ*

ATRI RUDRA[†]

*Department of Computer Science

Stanford University

{albertgu, rohanp, chrismre}@stanford.edu

[†]Department of Computer Science and Engineering

University at Buffalo, SUNY

atri@buffalo.edu

Abstract

Matrix-vector multiplication is one of the most fundamental computing primitives that has been studied extensively. Given a matrix $\mathbf{A} \in \mathbb{F}^{N \times N}$ and a vector $\mathbf{b} \in \mathbb{F}^N$, it is known that in the worst-case $\Theta(N^2)$ operations over \mathbb{F} are needed to compute $\mathbf{A}\mathbf{b}$. Many classes of *structured dense* matrices have been investigated that can be represented with $O(N)$ parameters, and for which matrix-vector multiplication can be performed with a sub-quadratic number of operations. One such class of structured matrices that admit near-linear matrix-vector multiplication are the *orthogonal polynomial transforms* whose rows correspond to a family of orthogonal polynomials. Other well known classes include the Toeplitz, Hankel, Vandermonde, Cauchy matrices and their extensions (e.g. confluent Cauchy-like matrices) that are all special cases of a low *displacement rank* property. In this paper, we identify a notion of *recurrence width* t of matrices \mathbf{A} so that such matrices can be represented with $t^2 N$ elements from \mathbb{F} . For matrices with constant recurrence width we design algorithms to compute both $\mathbf{A}\mathbf{b}$ and $\mathbf{A}^T \mathbf{b}$ with a near-linear number of operations. This notion of width is finer than all the above classes of structured matrices and thus computes near-linear matrix-vector multiplication for all of them using the *same* core algorithm. We consider extensions and variants of this width to other notions that can also be reduced to the main algorithms. These reductions are captured through operations on certain types of Krylov matrices, which have been used in many previous works such as Lanczos' and Wiedemann's algorithms. Furthermore, we show that the subclasses of Krylov matrices we use also exhibit a simple recurrence width structure. Technically, our work unifies, generalizes, and (we think) simplifies existing state-of-the-art results in structured matrix-vector multiplication. Finally, we show how applications in areas such as multipoint evaluations of multivariate polynomials and computing linear sequences can be reduced to problems involving low recurrence width matrices.

1 Introduction

1.1 Background and Overview of Our Results

In this paper, we focus on matrices in $\mathbb{F}^{N \times N}$, where \mathbb{F} is any field, for which one can design a near-linear time matrix-vector multiplication algorithm: i.e. an algorithm that takes $O(N \log^{O(1)} N)$ (which we will denote by $\tilde{O}(N)$) operations over \mathbb{F} . Such algorithms have been dubbed *superfast* algorithms in the matrix-vector multiplication literature [39]. Since superfast algorithms do not have time to read all $\Theta(N^2)$ elements of dense matrices, matrices that admit superfast multiplication algorithms must be structured in the sense of being expressible with a small number of parameters. Many problems such as the Discrete Fourier Transform, polynomial and rational multipoint evaluation/interpolation, and orthogonal polynomial projections can be expressed as matrix-vector multiplication involving dense, structured matrices, and specialized superfast algorithms have been designed for many of these specific problems [17, 22, 38]. We are interested in unifying these problems and their notions of structure. To this end, we introduce the concept of *recurrence width*, a new measure of the structural complexity of certain dense matrices, and design simple and general superfast matrix-vector multiplication algorithms for low-width matrices.

We believe that our strongest contribution is in showing that these existing classes of matrices, which were seemingly disparate and handled by different specialized algorithms, all fall under the umbrella of low recurrence width and can be handled with one class of algorithms.

Perhaps the poster child for matrices that allow for superfast vector multiplication is the Discrete Fourier Transform. The famous Fast Fourier transform (or FFT) allows matrix-vector multiplication for the Fourier matrix in $O(N \log N)$ operations [15]. Since then numerous followup works have extended this algorithm to solve the matrix vector multiplication algorithm for other, increasingly general structured dense matrices. To motivate our notion of structure, we will focus on the two strands of work that inspired much of our work, and describe how we capture previous results in these areas.

Orthogonal polynomial transforms The first strand of work relates to *orthogonal polynomial transforms* [1–3, 14]. The transforms can be expressed as matrix-vector multiplication involving matrices \mathbf{A} containing a family of orthogonal polynomials evaluated at points.

Definition 1.1. Let $f_0(X), \dots, f_{N-1}(X)$ be a collection of polynomials over a field \mathbf{F} and z_0, \dots, z_{N-1} be a set of points. The discrete polynomial transform matrix \mathbf{A} is defined by $\mathbf{A}_{ij} = f_i(z_j)$. The discrete polynomial transform of a vector \mathbf{b} with respect to the f_i and z_j is given by the product $\mathbf{A}\mathbf{b}$.

When the f_i are a family of orthogonal polynomials, we are left with an orthogonal polynomial transform. Further, these polynomials satisfy the following three term recurrence:

$$f_{i+1}(X) = (a_i X + b_i) f_i(X) + c_i f_{i-1}(X), \quad (1)$$

where $a_i, b_i, c_i \in \mathbb{F}$. Driscoll, Healy and Rockmore present an algorithm to perform the orthogonal polynomial transform in $O(N \log^2 N)$ operations [17]. In our first main result, we extended this class of transforms to polynomials that satisfy a more general recurrence.

Definition 1.2. An $N \times N$ matrix \mathbf{A} has recurrence width t if the polynomials $f_i(X) = \sum_{j=0}^{N-1} \mathbf{A}[i, j] X^j$ satisfy

$$f_{i+1}(X) = \sum_{j=0}^t g_{i,j}(X) f_{i-j}(X), \quad (2)$$

where the polynomials $g_{i,j} \in \mathbb{F}[X]$ are of degree at most $j + 1$.

This is the most basic notion of recurrence width that can be kept in mind as a prototypical example. A more general definition that captures the other strands of work and applications is presented in Definition 2.1.

We note that the orthogonal polynomial recurrence implies that orthogonal polynomial transforms are essentially recurrence width 1 matrices (the exact connection is in Section 9.2). We show in Section 8 that our notion of recurrence width forms a strong hierarchy. In particular, we construct a simple matrix \mathbf{A} with recurrence width $t + 1$ such that no matrix with recurrence width t can approximate $\mathbf{A}\mathbf{b}$ for all \mathbf{b} to any reasonable accuracy. Consequently, our extension is a meaningful one; the class of matrices we handle cannot be captured by previous results on orthogonal polynomials. This also justifies our use of the term width in this setting.

The recursive structure allows us to generate matrix vector multiplication algorithms for both matrices \mathbf{A} and \mathbf{A}^T in a simple and general way. Promisingly, our algorithms are optimal in the sense of matching the size of its parametrization, up to poly-log factors.

Theorem 1.3. *Given matrix \mathbf{A} with recurrence width t and $\tilde{O}(t^\omega N)$ pre-processing operations (where ω is the matrix-matrix multiplication exponent), the products $\mathbf{A}^T \mathbf{b}$ and $\mathbf{A}\mathbf{b}$ can be computed in $\tilde{O}(t^2 N)$ for any vector \mathbf{b} .*

Both parts of Theorem 1.3 are proven in Theorems 3.5 and 5.3 respectively.

This multiplication complexity is equal to the worst-case input size of a matrix with recurrence width t (and if $\omega = 2$, so is the pre-processing time). For example, we recover the bounds of Driscoll et al. [17], and Section 9.2 demonstrates how their algorithm can actually be viewed as a direct application of ours. In certain cases including orthogonal polynomials we also get matrix vector multiplication for \mathbf{A}^{-1} and $(\mathbf{A}^{-1})^T$.

Displacement rank The second relevant strand of work are results for matrices with low *displacement rank*. The notion of displacement rank (which was defined in the seminal work of Kailath et al. [28]) is defined as follows. Given any pair of matrices (\mathbf{L}, \mathbf{R}) , the displacement rank of \mathbf{A} with respect to (\mathbf{L}, \mathbf{R}) is the rank of the *error matrix*:

$$\mathbf{E} = \mathbf{L}\mathbf{A} - \mathbf{A}\mathbf{R}. \quad (3)$$

To the best of our knowledge, the most powerful results on matrix-vector multiplication for matrices with low displacement rank are in the work of Olshevsky and Shokrollahi [38], who show that any matrix with a displacement rank of r with respect to Jordan form matrices \mathbf{L} and \mathbf{R} can be multiplied by an arbitrary vector with $\tilde{O}(rN)$ operations. (They use these results and matrices to solve the Nevalina-Pick problem as well as solve the interpolation step in some list decoding algorithms in a previous work [37].) Recall that Jordan normal form matrices are special cases of matrices where only the diagonal and superdiagonal can be non-zero. In other words, both \mathbf{L} and \mathbf{R} are *2-band upper triangular matrices*. In this work we show that when both \mathbf{L} and \mathbf{R} are triangular t -band matrices, any matrix with displacement rank of r with respect to such matrices admits fast multiplication.

Theorem 1.4. *Let \mathbf{L} and \mathbf{R} be triangular t -band matrices sharing no eigenvalues, and let \mathbf{A} be a matrix such that $\mathbf{L}\mathbf{A} - \mathbf{A}\mathbf{R}$ has rank r . Then \mathbf{A} and \mathbf{A}^T can be multiplied by any vector \mathbf{b} in $\tilde{O}(t^2(t + r)N)$ operations.*

Theorem 1.4 is proven formally in Corollary 9.2. Our results recover the work presented in Olshevsky and Shokrollahi [38] - that is, when \mathbf{L} and \mathbf{R} are Jordan form matrices sharing no eigenvalues - and we believe that our algorithm even for their setting is simpler and less specialized.¹ Furthermore, our result can handle more general matrices than band matrices; in general, if one of \mathbf{L} and \mathbf{R} is a band matrix and the other admits superfast multiplication for the associated Krylov matrix (Theorem 9.2).

We find this connection compelling because the set of matrices with low recurrence width and those with low displacement rank seem to be widely different. Indeed the existing algorithms for the class of orthogonal polynomials [17] and low displacement rank [38] look very different. Specifically, the algorithm of Driscoll et al. [17] is a divide and conquer algorithm, while that of Olshevsky and Shokrollahi [38] (and preceding works) heavily exploit structural algebraic properties on matrices with low displacement ranks. We believe that our strongest conceptual contribution is showing that both of these classes of matrices can be handled with *one* class of algorithms. In particular, it turns out that if one allows for *error polynomials* in recurrences for matrices with recurrence width t then this is enough to handle the case of low displacement rank. In Section 2 we present the more abstract recurrence

¹Olshevsky and Shokrollahi make the claim that more general results hold with respect to the eigenvalue condition, but we note that constraints need to be made on \mathbf{L} and \mathbf{R} . For example, every matrix \mathbf{A} has low displacement rank with respect to Jordan matrices $\mathbf{L} = \mathbf{R} = \mathbf{I}$, so there cannot be a general multiplication algorithm. See Section 9.1 for further discussion.

that captures both of these classes of matrices. More importantly, we present efficient algorithms that work for these general recurrences. We believe that unifying these existing threads of disparate work is interesting in its own right.

As we have pointed out earlier, orthogonal polynomials and low displacement rank matrices have applications in numerous areas from signal processing to machine learning. Indeed orthogonal polynomials have their own dedicated conference [4]. For more details on matrices with low displacement rank as well as its application, please refer to the survey by Kailath and Sayed [29]. Our matrices naturally inherit these applications.

1.2 Our Algorithm and Techniques

As in all algorithms on structured matrices, the main challenge is in manipulating the alternate compact representation of the matrix directly. Although the recurrence (2) is represented with $O(N)$ values (for a fixed t), the polynomials $f_i(X)$ it produces are large. Unlike conventional linear recurrences on scalars, the total output size of this recurrence is quadratic because the polynomial degrees grow linearly. The first hurdle we clear is compressing this output: we use a divide-and-conquer algorithm to consider the large polynomials $f_{N/2}(X), \dots, f_{N-1}(X)$ separately, and use the recurrence structure to extract out a uniform portion of each polynomial and reduce their degrees to about $N/2$.

A second challenge is in increasing the power and flexibility of the representation (2) to capture a larger class of matrices. We will allow the recurrence coefficients to be rational functions instead of polynomials, where they are interpreted in an appropriate quotient ring of polynomials. We also allow the recurrence polynomials to be evaluated at a matrix \mathbf{R} , which produces a different matrix \mathbf{A} than the one from reading off the coefficients.

These generalizations can be unified by factoring out terms that capture the evaluation of certain matrices at polynomials or rational functions. We represent these as *Krylov matrices*, which appear in independently interesting settings on numerical methods such as the Lanczos eigenvalue algorithm and Wiedemann's kernel vector algorithm [30, 42]. The problem of evaluating a function of a matrix is itself a well-studied and useful problem [26], and certain strong representations of the matrix (such as the SVD or eigendecomposition) often play a central role. In contrast, we show that the cases we are interested in can be represented in a computationally simpler way with recurrence width.

Basic algorithms By the recurrence, each $f_i(X)$ can be written as a linear combination of the $f_0(X), \dots, f_t(X)$, where the coefficients of the combination are polynomials of degree approximately i (this is formalized in Lemmas 2.6 and 2.7). However, for $i \geq N/2$, this can be broken down into two parts: the linear combination of $f_i(X)$ in terms of $f_{N/2}(X)$ through $f_{N/2+t}(X)$, and the dependence of $f_{N/2}(X), \dots, f_{N/2+t}(X)$ on $f_0(X), \dots, f_t(X)$. The former has size approximately $i - N/2$, and the latter $N/2$. Thus the common theme is: break the recurrence into two recurrences of half the size, the first half initialized from 0, the second half initialized from $N/2$, and the 'jump' from $f_0(X), \dots, f_t(X)$ to $f_{N/2}(X), \dots, f_{N/2+t}(X)$.

We can illustrate the basic algorithms for $\mathbf{A}\mathbf{b}$ and $\mathbf{A}^T\mathbf{b}$ with a toy example consisting of polynomials $f_i(X)$ satisfying equation (2) with $f_0(X) = 1$ and $t = 0$. In this case, the polynomials can be easily factored as $f_i(X) = \prod_{j=0}^{i-1} g_{j,0}(X)$. First, the product $\mathbf{A}^T\mathbf{b}$ is equivalent to computing the sum $\sum_{i=0}^{N-1} \mathbf{b}[i] f_i(X)$ because of the correspondence between vectors and polynomials. This sum can be broken into

$$\sum_{i=0}^{N-1} \mathbf{b}[i] f_i(X) = \left(\sum_{i=0}^{N/2-1} \mathbf{b}[i] g_{0,0}(X) \cdots g_{i-1,0}(X) \right) + f_{N/2}(X) \left(\sum_{i=N/2}^{N-1} \mathbf{b}[i] g_{N/2,0}(X) \cdots g_{i-1,0}(X) \right),$$

which are two smaller sums of the same form, provided we can compute $f_{N/2}(X) = g_{0,0}(X) \cdots g_{N/2-1,0}(X)$. Thus the entire sum can be computed in near-linear operations if we know $\prod_{i=kN/2^d}^{(k+1)N/2^d-1} g_{i,0}(X)$ for all $d \in [\log N]$, $k \in [2^d]$. These can be pre-computed because they depend only on \mathbf{A} and not \mathbf{b} .

Computing the product $\mathbf{A}\mathbf{b}$ is more difficult because each polynomial $f_i(X)$ is not treated as a 'single entity' as in $\mathbf{A}^T\mathbf{b}$. Rather, the coefficients are individually manipulated - in particular, we need to compute the dot product of the coefficient vector of $f_i(X)$ with \mathbf{b} . Here we use the observation that the *dot product* between two vectors is a specific element of their *convolution*. Thus if $b(X) = \sum \mathbf{b}[i] X^{N-i}$, the entries of $\mathbf{A}\mathbf{b}$ are the coefficients of X^{N-1} in

$b(X), g_{0,0}(X)b(X), \dots, g_{0,0}(X) \cdots g_{i,0}(X)b(X), \dots, g_{0,0}(X) \cdots g_{N-1,0}(X)b(X)$. The second half of this is the coefficients of X^{N-1} in $f_{N/2}(X)b(X), \dots, f_{N/2}(X)g_{N/2,0}(X) \cdots g_{N-1,0}(X)b(X)$ which becomes a similar problem of half the size by defining $b'(X) = f_{N/2}(X)b(X)$. The details of this algorithm are in Section 5.

Extensions of the Recurrence Many of our applications require considering more complicated recurrences than (2), and one of our main technical contributions is identifying and addressing these generalizations. First note that the matrix \mathbf{A} in Definition 1.2 is fully defined by any polynomials $\{f_0(X), \dots, f_t(X)\}$ and $\{g_{i,j}(X)\}$ (if the recurrence generates a $f_i(X)$ of degree more than $N-1$, cut off the higher order terms or in general reduce it modulo a degree N polynomial $M(X)$). Now if the recurrence coefficients $g_{i,j}(X)$ were allowed to be degree $O(N)$, the recurrence would have full representative power (that is, any matrix \mathbf{A} can be produced), but the parameterization would be too large. A compromise is that sometimes even if the polynomial $g_{i,j}(X)$ has high degree, it can be represented compactly in a different way. Our first extension is treating these recurrence coefficients as elements of a polynomial modulus represented as a rational fraction. We express this idea with the recurrence

$$D_{i+1}(X) \cdot f_{i+1}(X) \equiv \sum_{j=0}^t g_{i,j}(X) f_{i-j}(X) \pmod{M(X)} \quad (4)$$

where for each i , $\gcd(D_i(X), M(X)) = 1$ and we define $f_{i+1}(X)$ to be the unique polynomial of degree less than N that satisfies the above equation.

A second extension of equation (2) generalizes the way we were transforming a polynomial recurrence into a matrix. If a matrix \mathbf{A} satisfies Definition 1.2, then its rows $\mathbf{f}_i = \mathbf{A}[i, :]^T$ satisfy a relation $\mathbf{f}_i = \sum_{j=0}^t g_{i,j}(\mathbf{S}) \mathbf{f}_{i-j}$; in other words the recurrence generates vectors by evaluating at the shift matrix \mathbf{S} which is 1 on the subdiagonal.² Thus a natural generalization is replacing \mathbf{S} with an arbitrary matrix in the recurrence.

$$\mathbf{f}_{i+1} = \sum_{j=0}^t g_{i,j}(\mathbf{R}) \mathbf{f}_{i-j} \quad (5)$$

Our final main extension is allowing for the rows of \mathbf{A} to satisfy (2) with some error terms. The recurrence takes the form

$$f_{i+1}(X) = \sum_{j=0}^t g_{i,j}(X) f_{i-j}(X) + E_{i+1}(X) \quad (6)$$

where the matrix \mathbf{E} formed by putting the coefficients of $E_i(X)$ on its rows has low rank. The connection between recurrence width and rank is explained in Section 6.3.

These extensions can be combined in different ways to capture disparate applications. Our connection to multipoint evaluation of multivariate polynomials involves matrices that have the form of both (6) and (4), and the solution to (5) itself reduces to a problem of that form. Of particular interest is that in Section 9.1, we show how the displacement rank equation (3) (where \mathbf{L}, \mathbf{R} are triangular t -banded matrices) allows us to write the rows of \mathbf{A} in the form of a recurrence using all three of these extensions, which is solved in Theorem 6.6:

$$D_{i+1}(X) \cdot \mathbf{f}_{i+1} \equiv \left(\sum_{j=0}^t g_{i,j}(\mathbf{R}) \mathbf{f}_{i-j} \right) + \mathbf{E}[i+1, :] \pmod{c_{\mathbf{R}}(\mathbf{R})}, \quad (7)$$

where $c_{\mathbf{R}}(X)$ is the characteristic polynomial of \mathbf{R} .

The extensions in (4) and (6) can be reduced fairly directly to the basic recurrence (2); this is done in Section 6.1 and Section 6.3 respectively. Next, we elaborate on the extension in (5) because it involves an intermediate problem with potentially broader uses.

²The equivalence follows from the isomorphism between polynomials of degree less than N and triangular Toeplitz matrices.

Matrix Functions and Krylov Efficiency Recurrence (5) involves evaluating functions at a matrix \mathbf{R} ; we focus on triangular t -banded matrices which is sufficient for the displacement rank application. Classical ways of computing these matrix functions use natural decompositions of the matrix, such as its singular value/eigendecomposition, or Jordan normal form $\mathbf{R} = \mathbf{A}\mathbf{J}\mathbf{A}^{-1}$. The evaluation of an analytic function then becomes $f(\mathbf{R}) = \mathbf{A}f(\mathbf{J})\mathbf{A}^{-1}$, which admits superfast multiplication if each component does. We note that the Jordan decomposition is generally hard to compute and much work has been done on computing specific matrix functions without going through the Jordan form [41]. Despite this, in Section 9.5 we show that it is possible to compute the Jordan decomposition quickly for several special subclasses of the matrices we are interested in, using techniques involving low-width recurrences.

However, solving (5) has more structure and is easier than evaluating general matrix functions. Consider again the simplified case when $t = 0$. Fixing the $g_{i,j}(X)$, we can define the polynomials $f_i(X)$ by (2). Note that $\mathbf{f}_i = g_{i-1,0}(\mathbf{R}) \cdots g_{0,0}(\mathbf{R})\mathbf{f}_0 = f_i(\mathbf{R})\mathbf{f}_0$. So we can factor $\mathbf{A} = \mathbf{A}'\mathbf{K}$, where \mathbf{A}' is the coefficient matrix of the basic polynomial recurrence (Definition 1.2), and \mathbf{K} is the matrix whose i -th column is $\mathbf{R}^i\mathbf{f}_0$. Thus this reduces to the basic recurrence (2) if we can multiply by \mathbf{K} , the *Krylov matrix* on \mathbf{R} and \mathbf{f}_0 .³ We say that \mathbf{R} is *Krylov efficient* if all of its Krylov matrices admit superfast multiplication.

In Section 7, we show that the Krylov matrix itself satisfies a recurrence of type (7) of recurrence width t . Using our established results, this gives a single $\tilde{O}(t^2N)$ algorithm that unifies the previous subclasses with Jordan decompositions, and implies all triangular banded matrices are Krylov efficient.

We remark that the Krylov efficiency concept does not apply only to our problem. If \mathbf{K} is the Krylov matrix on \mathbf{A} and \mathbf{b} , then $\mathbf{K}\mathbf{b} = \sum \mathbf{b}[i]\mathbf{A}^i\mathbf{x}$ is naturally related to contexts involving Krylov subspaces, matrix polynomials, and so on. The product $\mathbf{K}^T\mathbf{b} = [\mathbf{b} \cdot \mathbf{x}, \mathbf{b} \cdot \mathbf{A}\mathbf{x}, \mathbf{b} \cdot \mathbf{A}^2\mathbf{x}, \dots]$ is also useful; it is the first step in the Wiedemann algorithm for computing the minimal polynomial or kernel vectors of a matrix \mathbf{A} [30].

Quick comparison with existing work Finally, we again highlight that we consider our techniques to be relatively simple, but they work just as well as more involved techniques that only work in specific cases. In many of the past cases, the algorithms for $\mathbf{A}\mathbf{b}$ and $\mathbf{A}^T\mathbf{b}$ would require different techniques, but in our case we get both the results for a larger class of structured dense matrices with essentially the same idea. Further, some of the existing results that are based in structural results (e.g. those based on the Bezoutian [22] and on displacement rank [38]) seems to invoke some fairly sophisticated algebraic techniques while our results use mostly combinatorial techniques (in particular divide and conquer) except for polynomial multiplication and division. Because of their simplicity, we believe that our algorithms are also practical and should be competitive against existing algorithms that have been implemented. Our preliminary experimental results, presented in Section 9.7, have been encouraging and we plan to explore more rigorous experimental results in future work. We present a more detailed comparison with related work in Appendix A.

1.3 Some Consequences and More Context

We believe that given the fundamental nature of the matrix vector multiplication problem, our main contribution is presenting the notion of recurrence width for matrices for which we give optimal algorithms (up to poly-log factors) and in the process unify quite a few known results. However, our results have consequences beyond just the results in matrix-vector multiplication. We collect some of our favorite ones here.

A natural question to ask is if the matrices with recurrence width t (for say non-constant t) contain any interesting class of matrices beyond the fact that (i) these contain both the classes of matrices considered in [17] and [38] and (ii) these matrices have nice algorithmic properties. We present a few examples of matrices that have been studied before that indeed fall in this general category of matrices. In particular, we present applications in the following order: (1) We present matrices that arise naturally in coding theory and (2) We present matrices that arise when computing some well known sequences. Further, we would like to point out that our generic algorithms solve problems from these different areas each of which typically have their own home-grown algorithms and in some cases (e.g. Bernoulli numbers) we almost match the best known runtimes of these more specific algorithms.

³The image of this matrix is the Krylov subspace on \mathbf{R} and \mathbf{f}_0 .

Matrices from Coding Theory. We first observe that the problem of multipoint evaluation of multivariate polynomials (i.e. given a multivariate polynomial $f(X_1, \dots, X_m)$ over \mathbb{F} and N evaluation points $\mathbf{a}_1, \dots, \mathbf{a}_N \in \mathbb{F}^m$ output the vector $(f(\mathbf{a}_i))_{i \in [N]}$) corresponds to matrix-vector multiplication for matrices with recurrence width t (where t is polynomial but sub-linear in N). Further, the encoding problem of recently introduced multiplicity codes (which have excellent local decoding properties [32–34]), which in turn corresponds to evaluating multivariate polynomials and all of their derivatives up to a certain order, again corresponds to matrix vector multiplication with a matrix with recurrence width t (where t is polynomial but sub-linear in N). This was already observed in the context of list decoding Reed-Solomon codes for the case of $m = 2$ by Olshevsky and Shokrollahi [37] (we extend this observation to $m > 2$). While we do not prove any new upper bounds for these problems, it turns out that the connection to multipoint evaluation of multivariate polynomials has some interesting complexity implications for our result. In particular, recall that the worst-case input size of a matrix with recurrence width t is $\Theta(t^2 N)$ and our algorithms are optimal with respect to this input measure (assuming $\omega = 2$). However, it is natural to wonder if one can have faster algorithms with respect to a more per-instance input size. A natural such case would be to represent the coefficients $g_{i,j}(X)$ in (2) as *sparse* polynomials and count the input size as the total number of non-zero coefficients in all the $g_{i,j}(X)$'s. Another natural representation would be for cases where we can guarantee $\deg(g_{i,j}(X)) \leq d < t$: in such a case can we design efficient algorithms with respect to the input size $\Theta(tdN)$? In both cases we show that improving our generic algorithm substantially will improve the state of the art results in multipoint evaluation of multivariate polynomials over arbitrary fields.⁴ We present the details in Section 10.

Computing sequences. Second we observe that the recurrences that we consider for our matrices actually have *holonomic* sequences as a special case, which are some of the very well studied sequences and many algorithms for such sequences have been implemented in algebra packages [46]. Additionally, the computation of many well-known sequences of numbers, including Stirling numbers of the second kind and Bernoulli numbers, are also very well studied problems. There have been some recent ad-hoc algorithms to compute such numbers (e.g. the current best algorithm to compute the Bernoulli numbers appears in [25]). Despite these sequences not being holonomic, our general purpose algorithms can still be used to compute them. In Section 9.4, we show how to compute the first N Bernoulli numbers in $O(N \log^2 N)$ operations, which recovers the same algorithmic bit complexity (potentially with a log factor loss) as the algorithms that are specific to Bernoulli numbers.

2 Preliminaries

2.1 Notation

We will use \mathbb{F} to denote a field and use \mathbb{R} and \mathbb{C} to denote the field of real and complex numbers respectively.⁵ The set of polynomials over \mathbb{F} and set of rational functions over \mathbb{F} will be denoted by $\mathbb{F}[X]$ and $\mathbb{F}(X)$ respectively. For polynomials $p(X), q(X) \in \mathbb{F}[X]$, we use the notation $p(X) \equiv q(X) \pmod{M(X)}$ to indicate equivalence modulo $M(X)$, i.e. $M(X) | (p(X) - q(X))$, and $p(X) = q(X) \pmod{M(X)}$ to specify $p(X)$ as the unique element of this equivalence class with degree less than $\deg M(X)$. We will sometimes consider polynomials over multiple indeterminates; if $f(X_1, \dots, X_k) \in \mathbb{F}[X_1, \dots, X_k]$, we use $\deg(f)$ to mean the maximum degree of its monomials (i.e. sum of powers of all indeterminates in a term), and $\deg_{X_i}(f)$ to denote the degree of f when viewed as a polynomial over X_i . For any integer $m \geq 1$, we will use $[m]$ to denote the set $\{1, \dots, m\}$. Unless specified otherwise, indices in the paper start from 0.

In this paper, vectors are boldset like \mathbf{x} and are column vectors unless specified otherwise. We will denote the i th element in \mathbf{x} by $\mathbf{x}[i]$ and the vector between the positions $[\ell, r] : \ell \leq r$ by $\mathbf{x}[\ell : r]$. For any subset $T \subseteq [N]$, \mathbf{e}_T denotes the characteristic vector of T . We will shorten $\mathbf{e}_{\{i\}}$ by \mathbf{e}_i .

Matrices will be boldset like \mathbf{M} and by default $\mathbf{M} \in \mathbb{F}^{N \times N}$. We will denote the element in the i th row and j th column by $\mathbf{M}[i, j]$. $\mathbf{M}[\ell_1 : r_1, \ell_2 : r_2]$ denotes the sub-matrix $\{\mathbf{M}[i, j]\}_{\ell_1 \leq i < r_1, \ell_2 \leq j < r_2}$. In particular we will use $\mathbf{M}[i, :]$

⁴We note that for *finite* fields, Kedlaya and Umans [31] have essentially solved this problem. The case for general fields however, is much less explored and (somewhat to our surprise) widely open.

⁵It can be assumed that \mathbb{F} is \mathbb{R} or \mathbb{C} because our applications only use these fields, but all results can be extended to general fields \mathbb{F} and also commutative rings. We also assume that \mathbb{F} supports the FFT; otherwise, the runtime bounds incur an extra $\log \log N$ factor.

and $\mathbf{M}[i, j]$ to denote the i th row and j th column of \mathbf{M} respectively. ($\mathbf{M}[0, 0]$ denotes the ‘top-left’ element of \mathbf{M} .) We will use \mathbf{S} to denote the shift matrix (i.e. $\mathbf{S}[i, j] = 1$ if $i = j + 1$ and 0 otherwise) and \mathbf{I} to denote the identity matrix. We will use \mathbf{V} to denote Vandermonde matrices, defined in Definition A.1. Given a matrix \mathbf{A} , we denote its transpose and inverse (assuming it exists) by \mathbf{A}^T (so that $\mathbf{A}^T[i, j] = \mathbf{A}[j, i]$) and \mathbf{A}^{-1} (so that $\mathbf{A} \cdot \mathbf{A}^{-1} = \mathbf{I}$). We will denote $(\mathbf{A}^T)^{-1}$ by \mathbf{A}^{-T} .

Given a matrix $\mathbf{M} \in \mathbb{F}^{N \times N}$ and a vector $\mathbf{b} \in \mathbb{F}^N$, the *Krylov matrix of \mathbf{M} generated by \mathbf{b}* (denoted by $\mathcal{K}(\mathbf{M}, \mathbf{b})$) is the $N \times N$ matrix whose i th column for $0 \leq i < N$ is given by $\mathbf{M}^i \cdot \mathbf{b}$. We say that \mathbf{M} is (α, β) -Krylov efficient if for every $\mathbf{b} \in \mathbb{F}^N$, we have that $\mathbf{K} = \mathcal{K}(\mathbf{M}, \mathbf{b})$ admits the operations $\mathbf{K}\mathbf{x}$ and $\mathbf{K}^T\mathbf{x}$ (for any $\mathbf{x} \in \mathbb{F}^N$) with $\tilde{O}(\beta N)$ many operations (with $\tilde{O}(\alpha N)$ pre-processing operations). (The $\tilde{O}(\cdot)$ notation is defined below.) Section 7 has examples of Krylov efficient matrices.

Finally, we address some issues related to runtime analyses. We will use $\tilde{O}(T(N))$ to denote $O(T(N) \cdot \log^{O(1)}(T(N)))$. We will denote the *size* of a polynomial $p(X) \in \mathbb{F}[X]$ as the degree of $p(X)$.

Throughout this paper, we assume some well-known results that are summarized in Appendix A.1.

2.2 Our Problem

In this paper, we are interested in the matrix-vector multiplication problem. In other words, given $\mathbf{A} \in \mathbb{F}^{N \times N}$ and $\mathbf{b} \in \mathbb{F}^N$, we want to compute $\mathbf{c} = \mathbf{A}\mathbf{b}$ such that for every $0 \leq i < N$:

$$\mathbf{c}[i] = \sum_{j=0}^{N-1} \mathbf{A}[i, j] \cdot \mathbf{b}[j].$$

For the rest of the paper we will assume that $N = 2^m$: this does not change the asymptotics but makes some of the subsequent notations simpler. We will be interested in matrices \mathbf{A} for which we can compute $\mathbf{A}\mathbf{b}$ and $\mathbf{A}^T\mathbf{b}$ efficiently. In particular, we will be interested in structured matrices \mathbf{A} such that its rows satisfy certain recurrences. For notational convenience, define $\mathbf{f}_i^T = \mathbf{A}[i, :]$.

Definition 2.1. Let \mathcal{R} be a ring and $\otimes : \mathcal{R} \times \mathbb{F}^N \rightarrow \mathbb{F}^N$ be an operator satisfying⁶

$$a \otimes (b \otimes \mathbf{z}) = (a \cdot b) \otimes \mathbf{z} \tag{8}$$

$$a \otimes \mathbf{z} + b \otimes \mathbf{z} = (a + b) \otimes \mathbf{z} \tag{9}$$

$$1 \otimes \mathbf{z} = \mathbf{z} \tag{10}$$

A matrix $\mathbf{A} \in \mathbb{F}^{N \times N}$ has recurrence width t if and only if its rows $\mathbf{f}_i^T = \mathbf{A}[i, :]$ satisfy

$$\mathbf{f}_{i+1} = \sum_{j=0}^t g_{i,j} \otimes \mathbf{f}_{i-j} \tag{11}$$

for $i > t$.

The \otimes operator provides an abstraction so that the standard form (11) captures the recurrence variants and extensions we are interested in. We note that the \otimes operator essentially induces a left \mathcal{R} -module structure over \mathbb{F}^N [18]. We focus on when \mathcal{R} is a subring or quotient ring of $\mathbb{F}(X)$ so that each $g_{i,j}$ can be represented by a rational fraction $g_{i,j}(X) \in \mathbb{F}(X)$.

The basic polynomial recurrence as in Definition 1.2 as well as its extensions (aside from the recurrence with error (6)) can all be viewed as various instantiations of the \otimes operator.

Definition 2.2. Let $\mathcal{R} = \mathbb{F}[X]$ and define $a(X) \otimes \mathbf{z}$ to be the convolution between the coefficient vector of $a(X)$ and \mathbf{z} .

⁶The definition can be extended naturally to $\mathbf{A} \in \mathcal{R}^{p \times q}$ and $\mathbf{Z} \in \mathbb{F}^{q \times r}$ where $\mathbf{A} \otimes \mathbf{Z}$ is defined through the usual matrix multiplication but replacing product with \otimes . This also satisfies the same properties, i.e. $\mathbf{A} \otimes (\mathbf{B} \otimes \mathbf{Z}) = (\mathbf{A}\mathbf{B}) \otimes \mathbf{Z}$, $\mathbf{A} \otimes \mathbf{Z} + \mathbf{B} \otimes \mathbf{Z} = (\mathbf{A} + \mathbf{B}) \otimes \mathbf{Z}$ and $\mathbf{I} \otimes \mathbf{Z} = \mathbf{Z}$.

This defines recurrence (2) and is the main type of recurrence we consider. The classic case is when $\deg(f_i) \leq i$ for $0 \leq i \leq t$, so that $\deg(g_{i,j}) \leq j+1$ implies that $\deg(f_i) \leq i$ for all i (which is true for orthogonal polynomials with $t=1$). When the former does not hold, it may happen that $\deg(f_i) \geq N$, but we can still define the matrix \mathbf{A} by cutting off the higher order terms, i.e. let $\mathbf{A}[i, j]$ be the coefficient of X^j in $f_i(X)$. This is more precisely an instance of the next definition of \otimes with $M(X) = X^N$.

Definition 2.3. Let $\mathcal{R} = \mathbb{F}[X]/(M(X))$ for $M(X) \in \mathbb{F}[X]$ of degree N , and define $a(X) \otimes \mathbf{z}$ as in Definition 2.2. More precisely, it is the coefficient vector of $b(X)$ where $b(X) = a(X) \cdot (\sum_{i=0}^{N-1} \mathbf{z}[i] X^i) \pmod{M(X)}$.

Note that Definition 2.3 captures recurrence (4): that equation is equivalent to writing $\mathbf{f}_i = \sum (g_{i,j}(X)/D_{i+1}(X)) \otimes \mathbf{f}_{i-j}$ using this definition of \otimes .

Definition 2.4. Let $\mathbf{R} \in \mathbb{F}^{N \times N}$ and $\mathcal{R} = \mathbb{F}[X]$ or $\mathcal{R} = \mathbb{F}[X]/(c_{\mathbf{R}}(X))$. Given $a(X) \in \mathcal{R}$, let $a(X) \otimes \mathbf{z} = a(\mathbf{R})\mathbf{z}$.

This captures recurrence (5). In Appendix B we show how Definition 2.4 captures another natural instantiation of \otimes involving matrices.

As usual the prototypical *polynomial recurrence* is through Definition 2.2. The other cases will reduce to this case: Definition 2.3 is covered in Section 6.1 and Definition 2.4 in Section 6.2. We call a *modular recurrence* one that is defined by Definition 2.3 and a *\mathbf{R} -matrix recurrence* one that is defined by Definition 2.4. Appendix B contains examples of matrices with low recurrence width under these definitions.

No matter which definition is used, we need to assume degree constraints on the recurrence coefficients in order to impose structure; otherwise, equation (11) can define any matrix.

Definition 2.5. We say that a recurrence satisfying (11) is (d, \bar{d}) -nice if each $g_{i,j}(X)$ can be represented with a fraction $p_{i,j}(X)/q_{i,j}(X) \in \mathbb{F}(X)$ such that

1. For each i , $q_{i,j}(X) = q_i(X)$ is equal for all j , and $\deg(q_i(X)) \leq \bar{d}$.
2. For all i, j , $\deg(p_{i,j}(X)) \leq d(j+1) + \bar{d}$.

Note that \bar{d} will be 0 in the settings of Definition 2.2 or Definition 2.4 when $\mathcal{R} = \mathbb{F}[X]$. If (d, \bar{d}) is not specified or clear from context, it is assumed to be $(1, 0)$; this corresponds to the basic polynomial recurrence case.

Finally, we will use $\mathbf{G} \in (\mathbb{F}(X))^{N \times (t+1)}$ to compactly represent the input: we will set $\mathbf{G}[i, j] = g_{i,j}(X)$ for $0 \leq i < N$ and $0 \leq j \leq t$. Further, we will use $\mathbf{F} \in (\mathbb{F}^N)^{t+1}$ to contain the initial condition, i.e. $\mathbf{F}[i] = \mathbf{f}_i$ for $0 \leq i \leq t$. A common instantiation of this vector will be when $\mathbf{f}_i = \delta_{i,j}$ for some $0 \leq j \leq t$. Note that the pair (\mathbf{G}, \mathbf{F}) completely specifies the recurrence and we will refer to it simply as (\mathbf{G}, \mathbf{F}) -recurrence.

2.3 The Work-Horse Lemma

Next, we present a lemma that formalizes the simple notion that one can re-initialize the recurrence from any $\mathbf{f}_k, \mathbf{f}_{k+1}, \dots, \mathbf{f}_{k+t}$ (where k need not be 0). This notation will be useful throughout the main algorithms.

First, for any $0 \leq i < n$, we will define the $(t+1) \times (t+1)$ matrix:

$$\mathbf{T}_i = \begin{pmatrix} g_{i,0}(X) & g_{i,1}(X) & \cdots & g_{i,t-1}(X) & g_{i,t}(X) \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \end{pmatrix}.$$

The above matrix in the literature is also called the *companion matrix*. And for any $\ell \leq r$, define

$$\mathbf{T}_{[\ell:r]} = \mathbf{T}_{r-1} \times \cdots \times \mathbf{T}_{\ell}.$$

Conceptually, the lemma below states that every term in the recurrence can be written as a combination of any $t+1$ consecutive terms and provides properties about these combination coefficients. It is helpful to consider the polynomial recurrence case (2), whence every term is a linear combination (with polynomial coefficients) of previous consecutive terms, as a prototypical example.

Lemma 2.6. Given a sequence $\mathbf{f}_0, \dots, \mathbf{f}_{N-1}$ specified by (11), for every index k , there exist $h_{i,j}^{(k)}(X) \in \mathbb{F}(X)$ for $0 \leq i < N - k, 0 \leq j \leq t$ such that

1. $\mathbf{f}_{i+k} = \sum_{j=0}^t h_{i,j}^{(k)}(X) \otimes \mathbf{f}_{k+j}$
2. $h_{i+1,j}^{(k)}(X) = \sum_{\ell=0}^t g_{i+k,\ell}(X) h_{i-\ell,j}^{(k)}(X)$ for $i \geq t$ and $h_{i,j}(X) = \delta_{i,j}$ for $i \leq t$.

The following result about the sizes of entries in $\mathbf{T}_{[\ell:r]}$ will be useful later.

Lemma 2.7. Let the recurrence in (11) be (d, \bar{d}) -nice. Then for any $0 \leq \ell \leq r < N$, the matrix $\mathbf{T}_{[\ell:r]} = D_{\ell,r}(X)^{-1} \mathbf{T}'_{\ell,r}$ where $\deg(D_{\ell,r}) \leq \bar{d}(r - \ell)$ and for all $0 \leq i, j \leq t$,

$$\deg(\mathbf{T}'_{[\ell:r]}[i, j]) \leq \bar{d}(r - \ell) + d \max((r - \ell + j - i), 0).$$

The proofs of these lemmas are deferred to Appendix C.

3 Computing $\mathbf{A}^T \mathbf{b}$

We consider the problem of computing $\mathbf{A}^T \mathbf{b}$ or equivalently $\mathbf{b}^T \mathbf{A}$:

$$\mathbf{c} = \sum_{i=0}^N \mathbf{b}[i] \mathbf{f}_i$$

for the standard recurrence satisfying (11) and Definition 2.2. We assume it is $(1, 0)$ -nice, i.e. $\deg(g_{i,j}(X)) \leq j + 1$.

We will solve this problem by isolating the effect of the \otimes operator and reducing the problem to a main computation just over $\mathbb{F}[X]$. Let $\mathbf{F} = [\mathbf{f}_t \cdots \mathbf{f}_0]^T$. Lemma 2.6 implies that \mathbf{f}_i is the last element of $\mathbf{T}_{[t:i+t]} \otimes \mathbf{F}$. Thus \mathbf{c} is the last element of

$$\sum_{i=0}^{N-1} \mathbf{b}[i] (\mathbf{T}_{[t:i+t]} \otimes \mathbf{F}) \quad (12)$$

This can be rewritten as $(\sum_{i=0}^{N-1} \mathbf{b}[i] \mathbf{T}_{[t:i+t]}) \otimes \mathbf{F}$. Thus it suffices to compute the last row of this sum in parenthesis. It can be decomposed as

$$\sum_{i=0}^{N/2-1} \mathbf{b}[i] \mathbf{T}_{[t:i+t]} + \left(\sum_{i=N/2}^{N-1} \mathbf{b}[i] \mathbf{T}_{[N/2+t:i+t]} \right) \mathbf{T}_{[t:N/2+t]}$$

Note that both sums are the same problem but of size $N/2$, corresponding to both halves of the recurrence. This motivates defining $\mathbf{P}_{\ell,r}$ to be the last row of $\sum_{i=\ell}^{r-1} \mathbf{b}[i] \mathbf{T}_{[\ell+t:i+t]}$, so that the desired answer is simply $\mathbf{P}_{0,N} \otimes \mathbf{F}$. Furthermore, this quantity satisfies the relation

$$\mathbf{P}_{\ell,r} = \mathbf{P}_{\ell,m} + \mathbf{P}_{m,r} \mathbf{T}_{[\ell:m]} \quad (13)$$

for any $\ell \leq m < r$, and thus can be computed with two recursive calls and a vector-matrix multiply over $\mathcal{R}(X)^{t+1}$.

Also, note that the $\mathbf{T}_{[\ell:r]}$ are independent of \mathbf{b} and the relevant ones will be pre-computed.

Finally, observe that when $r - \ell \leq t$, the last row of $\mathbf{T}_{[\ell:r]}$ is simply an indicator vector with a 1 in the $r - \ell$ -th spot from the end (conceptually, the companion matrices act as a shift before reaching the recurrence width). Thus we stop the recurrence when the problem size gets below $t + 1$.

We present the full details in Algorithm 1. The initial call is `TRANSPOSEMULT($\mathbf{b}, \mathbf{G}, m, t$)` (recall that $N = 2^m$) and the relevant $\mathbf{T}_{[\ell:r]}$ are assumed to be precomputed, which we will describe next.

The above discussion implies that Algorithm 1 is correct:

Theorem 3.1. Assuming all the required $\mathbf{T}_{[\cdot,\cdot]}$ are correctly computed, `TRANSPOSEMULT($\mathbf{b}, \mathbf{G}, m, 0$)` returns $\mathbf{P}_{0,N}$.

Algorithm 1 TRANSPOSEMULT

Input: $\mathbf{T}_{[bN/2^d : bN/2^d + N/2^{d+1}]}$ for $0 \leq d < m, 0 \leq b < 2^d$

Input: \mathbf{b}, a, k

Output: $\mathbf{P}_{k, k+2^a}$ = Last row of $\sum_{i=0}^{2^a-1} \mathbf{b}[k+i] \mathbf{T}_{[k:k+i]}$

1: **If** $2^a \leq t$ **then**

▷ Base case

2: **Return** $[0 \quad \dots \quad \mathbf{b}[k+2^a-1] \quad \dots \quad \mathbf{b}[k]]$

3: **else**

4: **Return** TRANSPOSEMULT($\mathbf{b}, a-1, k$) + TRANSPOSEMULT($\mathbf{b}, a-1, k+2^{a-1}$) $\mathbf{T}_{[k:k+2^{a-1}]}$

3.1 Pre-processing time

In this section, we will see how we can efficiently compute the matrices of polynomials $\mathbf{T}_{[bN/2^d : bN/2^d + N/2^{d+1}]}$ for $0 \leq d < m, 0 \leq b < 2^d$. Since we have assumed that N is a power of 2, these ranges can be expressed in terms of dyadic strings; we need to pre-compute $\mathbf{T}_{[s]}$ for all strings $s \in \{0, 1\}^*, |s| \leq \lg N$ (where we interpret $[s]$ as the corresponding dyadic interval in $[0, N-1]$). All the required matrices can be computed in a natural bottom-up fashion:

Lemma 3.2. *We can pre-compute $\mathbf{T}_{[s]}$ for all strings $s \in \{0, 1\}^*, |s| \leq \lg N$ with $O(N t^{\omega_{\mathbb{F}}} \log^2 N)$ operations, where two $n \times n$ matrices over \mathbb{F} can be multiplied with $O(n^{\omega_{\mathbb{F}}})$ many operations over \mathbb{F} .*

Proof. Fix an arbitrary s^* of length $\ell < \lg N$. We can compute $\mathbf{T}_{[s^*]} = \mathbf{T}_{[s^*0]} \cdot \mathbf{T}_{[s^*1]}$. Using the matrix multiplication algorithm, we have $O(t^{\omega_{\mathbb{F}}})$ polynomial multiplications to compute, where the polynomials are of size at most $\frac{N}{2^\ell} + t$ by Lemma 2.7. So computing $\mathbf{T}_{[s^*]}$ takes $O((N/2^\ell + t) t^{\omega_{\mathbb{F}}} \log N)$ operations. Computing $\mathbf{T}_{[s]}$ for all $|s| = \ell$ takes $O((N + t 2^\ell) t^{\omega_{\mathbb{F}}} \log N)$, and computing all $\mathbf{T}_{[s]}$ takes $O(N t^{\omega_{\mathbb{F}}} \log^2 N)$, as desired. \square

The above implies that

Corollary 3.3. *Pre-processing for Algorithm 1 can be done with $O(t^{\omega_{\mathbb{F}}} N \log^2 N)$ many operations over \mathbb{F} .*

3.2 Runtime analysis

We now analyze the runtime of Algorithm 1 assuming that the pre-processing step (i.e. computing all the required $\mathbf{T}_{[s]}$) is already done. We do a simple recursive runtime analysis.

Lemma 3.4. *After pre-processing, Algorithm 1 needs $O(t^2 N \log^2 N)$ operations over \mathbb{F} .*

Proof. Let $T(N)$ denote the number of operations needed in the call TRANSPOSEMULT($\mathbf{b}, \mathbf{G}, m, 0$). Then it is easy to see that

$$T(N) = 2T(N/2) + O(\tau(N)),$$

where $\tau(N)$ is the number of operations taken in Step 4. We will show that $\tau(N) = O((t^2 N + t^3) \log N)$. Finally, note that $\tau(t+1) = O(t)$. This immediately implies the claim (note that the $O(t^3)$ term ultimately gets multiplied by N/t since we stop the recursion when the input size is $t+1$).

To see why $\tau(N) = O((t^2 N + t^3) \log N)$, we first note that by Lemma 2.7, each element of vector $\mathbf{P}_{\ell, r}$ is a polynomial of size at most $r - \ell + t$. Thus the vector-matrix multiplication in Step 4 performs $O(t^2)$ multiplications of polynomials of size at most $N/2 + t$. Thus, the total number of operations used by the call to TRANSPOSEMULT($\mathbf{b}, \mathbf{G}, m, 0$) is $O(t^2(N + t) \log N)$, as desired. \square

⁷In the last estimate we note that when $|s^*| \leq \log_2(t+1)$, then all operations are over \mathbb{F} and no polynomials are involved. Thus the $t^{\omega_{\mathbb{F}}+1} 2^\ell$ term only gets added up to $\ell = \log(N/t)$ and thus we do not pay an extra factor of t in the final analysis.

3.3 Post-processing

We are not quite done yet with regard to computing $\mathbf{A}^T \mathbf{b}$ since after Algorithm 1, we still have to perform the step $\mathbf{P}_{0,N} \otimes \mathbf{F}$. Note that in the setting of Definition 2.2 this is just t multiplications over $\mathbb{F}[X]$ of polynomials of size at most N : this can be done with $O(tN \log N)$ operations. These observations along with Corollary 3.3 and Lemma 3.4 imply the following result:

Theorem 3.5. *For any matrix \mathbf{A} satisfying a $(1, 0)$ -nice recurrence (11) under Definition 2.2, with $O(t^{\omega_{\mathbb{F}}} N \log^2 N)$ pre-processing operations, the product $\mathbf{A}^T \mathbf{b}$ can be computed for any \mathbf{b} with $O(t^2 N \log^2 N)$ operations over \mathbb{F} .*

We make a final remark on an optimization that can be done in the matrix-matrix multiplication setting. First, note that (12) can be written as

$$\sum_{i=0}^{N-1} \begin{bmatrix} 0 & \cdots & 0 & \mathbf{b}[i] \end{bmatrix} (\mathbf{T}_{[t:i+t]} \otimes \mathbf{F})$$

Now suppose we are given r vectors $\mathbf{b}_1, \dots, \mathbf{b}_r$ and want to compute $\mathbf{A}\mathbf{b}_1, \dots, \mathbf{A}\mathbf{b}_r$. Defining $\mathbf{B}_i = \begin{bmatrix} 0 & \cdots & 0 & \mathbf{b}_i[i] \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & \mathbf{b}_i[i] \end{bmatrix}$,

it suffices to compute

$$\sum_{i=0}^{N-1} \mathbf{B}_i \mathbf{T}_{[t:i+t]} \quad (14)$$

We can call the above $\mathbf{P}_{0,N}$, defined the same way as before, and the recursive identity (13) still holds so we can still run Algorithm 1. The analysis in Lemma 3.4 showed that the bottleneck of the algorithm, which led to the t^2 runtime coefficient, is in the multiplication of a $1 \times t$ and $t \times t$ matrix. Thus for this batch query, the runtime will be $\tilde{O}(t_r^\alpha N)$ where t_r^α is the number of operations to multiply a $r \times t$ by $t \times t$ matrix.

Note in particular that $r = t$ multiplications can be computed simultaneously in $\tilde{O}(t^\omega N)$ time, and in general $t_r^\alpha \leq \min(rt^2, (1 + r/t)t^\omega)$. For large numbers of queries, the amortized multiplication time can be considered to be $\tilde{O}(t^{\omega-1} N)$. This shows

Corollary 3.6. *For any matrix \mathbf{A} satisfying a $(1, 0)$ -nice recurrence (11) under Definition 2.2, with $O((t+r)t^{\omega_{\mathbb{F}}-1} N \log^2 N)$ pre-processing operations, the product $\mathbf{A}^T \mathbf{B}$ can be computed for any $\mathbf{B} \in \mathbb{F}^{N \times r}$ with $O((t^2 + tr)N \log^2 N)$ operations over \mathbb{F} .*

4 Overview of the Rest of the Paper

We present the rest of our main matrix-vector multiplication algorithms in Sections 5 and 6. In particular, we present the algorithm to compute $\mathbf{A}\mathbf{b}$ for the simplest setting in Section 5, and show that it has the same complexity guarantees as Theorem 3.5. In Section 6 we address generalizations of the recurrence. Section 6.1 shows how to reduce a modular recurrence (4) to a basic polynomial recurrence. In Section 6.2, we show that a matrix recurrence (5) can be reduced by factoring out Krylov matrices. Section 6.3 deals with the case when our recurrences have error. This is crucial for extending our results to work for low displacement rank matrices (and beyond). One solution is fairly simple: for a recurrence with the error matrix having rank r , we can reduce the problem with error to $r + 1$ instantiations of the problem with no errors and use our previous results as a black box. We also provide a more complicated approach with a reduced runtime. We also show in Theorem 10.4 that doing substantially better than this reduction would improve the state of the art algorithms for multipoint evaluation of multivariate polynomials.

In Section 7 we present examples of some matrices that are Krylov efficient, most importantly general triangular banded matrices. Notably, this depends on a previous result in Section 6.1.

In Section 8, we show that our notion of recurrence width satisfies a hierarchy in a very strong sense: in particular to represent a matrix with recurrence width $t + 1$ with a matrix of recurrence width t needs an error matrix of rank $\Omega(N/t)$.

In Section 9, we consider some special cases of our general framework. In Sections 9.1 and 9.2, we show why low displacement rank matrices fall in our framework and how our results in some sense are simpler than known results for orthogonal polynomial transforms. In Section 9.3, we show how to speed up our algorithms for the case when the $g_{i,j}(X)$ s only depend on j (i.e. the ‘transition’ matrix \mathbf{T}_i is independent of i). In Section 9.4 we show how our general purpose algorithm can be used to compute Bernoulli numbers. In Section 9.5, we utilize our algorithm to compute the Jordan decomposition of some classes of matrices. In Section 9.6, we outline why our algorithms have good bit complexity and in particular, why they are competitive with special purpose algorithms [25] to compute Bernoulli numbers even in the bit complexity measure. Section 9.7 presents some preliminary experimental results.

Finally, in Section 10 we present the connections of our framework to multipoint evaluation of multivariate polynomials as well as coding theory. The aim of this section is two-fold. First this section presents matrices that have been studied in other contexts that have a small recurrence width in our setup (but do not easily fit into the existing notions of width). Second, this section shows that if the efficiency of our algorithms can be improved in some natural directions, then those improvements will immediately imply improvements over the state of the art algorithms for multipoint evaluation of multivariate polynomials over arbitrary fields (and with arbitrary evaluation points).

5 Computing \mathbf{Ab}

In this section we will solve the problem of computing

$$\mathbf{c} = \mathbf{Ab}.$$

for a basic polynomial recurrence: the rows of \mathbf{A} satisfy recurrence (11) under Definition 2.2 and the $g_{i,j}(X)$ are $(1,0)$ -nice. For the purposes of later generalizing to modular recurrences, we will actually consider a minor generalization where the coefficients are a nice family of rational functions. Namely, consider the polynomials $f_i(X)$ defined by recurrence

$$D_i(X)f_{i+1}(X) = \sum_{j=0}^t n_{i,j}(X) \cdot f_{i-j}(X). \quad (15)$$

where $\deg(D_i(X)) = \bar{d}$ and $\deg(n_{i,j}(X)) \leq d(j+1) + \bar{d}$; that is, the recurrence coefficients $g_{i,j}(X) = n_{i,j}(X)/D_i(X)$ are (d, \bar{d}) -nice. Furthermore, assume the starting conditions $f_i(X) : 0 \leq i \leq t$ satisfy

$$\prod_{j=0}^{N-1} D_j(X) \mid f_i(X) \quad (16)$$

(this condition is only to ensure that the recurrence generates polynomials, which will be shown).

Assume that the $f_i(X)$ have degrees bounded by $\bar{N} = (d + \bar{d})N$, and we will consider the problem of computing \mathbf{Ab} where the $\mathbf{A}[i, j]$ is the coefficient of X^j in $f_i(X)$ (note that $\mathbf{A} \in \mathbb{F}^{N \times \bar{N}}$, $\mathbf{b} \in \mathbb{F}^{\bar{N}}$).

We again emphasize that the setting when $(d, \bar{d}) = (1, 0)$ corresponds to the basic polynomial recurrence and should be considered the prototypical example for this section. In this case $\bar{N} = N$, $D_i(X) = 1$ for all i , and the divisibility assumptions on the starting polynomials (16) are degenerate.

The main idea of the algorithm is to use the following observation to compute \mathbf{Ab} . For any vector $\mathbf{u} \in \mathbb{F}^N$, define the polynomial

$$\mathbf{u}(X) = \sum_{i=0}^{N-1} u_i \cdot X^i.$$

Also define

$$\mathbf{u}^R = \mathbf{J} \cdot \mathbf{u},$$

for the *reverse* vector, where \mathbf{J} is the ‘reverse identity’ matrix. Finally, define $\text{Coeff}_i(p(X))$ to be the coefficient of the term X^i in the polynomial $p(X)$. With the above notations we have

Lemma 5.1. For any vector $\mathbf{u}, \mathbf{v} \in \mathbb{F}^N$, we have $\langle \mathbf{u}, \mathbf{v} \rangle (\langle \mathbf{u}, \mathbf{v}^R \rangle) = \text{Coeff}_{N-1}(\mathbf{u}(X) \cdot \mathbf{v}^R(X))$ ($\mathbf{u}(X) \cdot \mathbf{v}(X)$ resp.).

Proof. The proof follows from noting that the coefficient of X^{N-1} in $\mathbf{u}(X) \cdot \mathbf{v}^R(X)$ is given by

$$\sum_{j=0}^{N-1} \mathbf{u}[j] \cdot \mathbf{v}^R[N-1-j] = \sum_{i=0}^{N-1} \mathbf{u}[i] \cdot \mathbf{v}[i] = \langle \mathbf{u}, \mathbf{v} \rangle,$$

where we used that fact that $(\mathbf{v}^R)^R = \mathbf{v}$. □

For notational convenience, define $D_{[i:j]}(X) = \prod_{k=i}^{j-1} D_k(X)$ (i, j can be out of the range $[t : N]$ with the convention $D_i(X) = 1$ for i outside the range). By Lemma 5.1, we know that $\mathbf{c}[i] = \text{Coeff}_{\tilde{N}-1}(f_i(X)\mathbf{b}^R(X))$. By Lemma 2.6, this means that $\mathbf{c}[i]$ is the coefficient of $X^{\tilde{N}-1}$ of

$$\mathbf{eT}_{[t:i+t]}D_{[t:N+t]}(X)\mathbf{F},$$

where $\mathbf{F} \in \mathbb{F}^{t+1}$ is defined by

$$\mathbf{F}[t-i] = \frac{f_i(X)\mathbf{b}^R(X)}{D_{[t:N+t]}(X)}, \quad (17)$$

and \mathbf{e} is the row vector $[0 \ \cdots \ 0 \ 1]$ (i.e. so we are considering the last row of $\mathbf{T}_{[t:i+t]}$). Note that this expression is well-defined because $\mathbf{T}_i D_i(X)$ is a matrix of polynomials, and \mathbf{F} is a vector of polynomials by (16), so the resulting product is a polynomial.

Note that the first half of these expressions can be written as $(\mathbf{eT}_{[t:i+t]}D_{[t:N/2+t]}(X))(D_{[N/2+t:N+t]}(X)\mathbf{F})$ for $0 \leq i < N/2$. By Lemma 2.7, the left term $\mathbf{T}_{[t:i+t]}D_{[t:N/2+t]}(X)$ is a matrix with polynomial entries where the last row has degrees bounded by $di + \tilde{d}N/2 < \tilde{N}/2$. So the coefficient of \tilde{N} in the whole product depends only on the higher-order $\tilde{N}/2$ coefficients of the right term $D_{[N/2+t:N+t]}(X)\mathbf{F}$. For convenience, define this operator which reduces a polynomial to one on its higher order coefficients

$$\text{Reduce}(f(X), n) = \frac{(f(X) \bmod X^n) - (f(X) \bmod X^{n/2})}{X^{n/2}}$$

Thus if we define $\mathbf{F}_\ell[i] = \text{Reduce}(D_{[N/2+t:N+t]}(X)\mathbf{F}[i], \tilde{N})$, then $\mathbf{c}[i]$ for $0 \leq i < N/2$ is the coefficient of $X^{\tilde{N}/2-1}$ in

$$\mathbf{eT}_{[t:i+t]}D_{[t:N/2+t]}(X)\mathbf{F}_\ell.$$

Note that this has the same form as the original problem, but with every term of half the size.

Similarly, we examine the second half of the answer. We want the coefficient of $X^{\tilde{N}-1}$ in $\mathbf{eT}_{[t:N/2+i+t]}D_{[t:N+t]}(X)\mathbf{F}$ for $0 \leq i < N/2$, which can be written as $(\mathbf{eT}_{[N/2+t:N/2+i+t]}D_{[N/2+t:N+t]}(X))(\mathbf{T}_{[t:N/2+t]}D_{[t:N/2+t]}(X)\mathbf{F})$. Note once again that the left matrix has degrees bounded by $\tilde{N}/2$, so we only need the higher order $\tilde{N}/2$ coefficients of the polynomials in the right term. Thus defining $\mathbf{F}_r[i] = \text{Reduce}(\mathbf{T}_{[t:N/2+t]}D_{[t:N/2+t]}(X)\mathbf{F}[i], \tilde{N})$, then $\mathbf{c}[N/2+i] : 0 \leq i < N/2$ is the coefficient of $X^{\tilde{N}/2-1}$ in

$$\mathbf{eT}_{[N/2+t:N/2+i+t]}D_{[N/2+t:N+t]}(X)\mathbf{F}_r$$

which is once again a problem of half the size of the original.

The algorithm is formalized in Algorithm 2, the correctness of which follows from the above discussion. The initial call is to $\text{MATRIXVECTMULT}(\mathbf{F}, m, 0)$, where we assume that $N = 2^m$.

We argue that Algorithm runs efficiently. In particular,

Lemma 5.2. A call to $\text{MATRIXVECTMULT}(\mathbf{F}, m, t)$ takes $O(t^2(\tilde{d} + d)N \log^2 N)$ many operations.

Proof. First we note that $D_{[bN/2^d : bN/2^d + N/2^{d-1}]}(X)$ can be computed for all $0 \leq d < m, 0 \leq b < 2^d$ in $O(\tilde{d}N \log^2 N)$ operations by a straightforward divide and conquer, so we compute these first within the time bound.

Also note that the base case is just t multiplications of polynomials of size $(d + \tilde{d})t$, which takes $O((d + \tilde{d})t^2 \log t)$ operations to compute.

Algorithm 2 MATRIXVECTMULT(\mathbf{F}, a, k)

Input: $\mathbf{T} \left[\frac{bN}{2^d} : \frac{bN}{2^d} + \frac{N}{2^{d+1}} \right]$ for $0 \leq d < m, 0 \leq b < 2^d$

Input: \mathbf{F}, a, k , such that $\deg(\mathbf{F}[i]) \leq (d + \bar{d})2^a$ for $0 \leq i \leq t$

Output: $\mathbf{c}[i] = \text{Coeff}_{X^{(d+\bar{d})2^{a-1}}}(\mathbf{eT}_{[k:k+i]}D_{[k:k+2^a]}(X)\mathbf{F})$, for $0 \leq i < 2^a$

1: $n \leftarrow 2^a$

2: **If** $n \leq t$ **then**

▷ Base case

3: $\mathbf{c}[k+i] \leftarrow \text{Coeff}_{X^{(d+\bar{d})n-1}}(D_{[k:k+2^a]}(X)\mathbf{F}[t-i])$ for $0 \leq i < n$

4: $\mathbf{F}_\ell \leftarrow \text{Reduce}(D_{[k+n/2:k+n]}(X)\mathbf{F}, (d + \bar{d})n)$

5: $\mathbf{c}[k : k + n/2] \leftarrow \text{MATRIXVECTMULT}(\mathbf{F}_\ell, a-1, k)$

6: $\mathbf{F}_r \leftarrow \text{Reduce}(\mathbf{T}_{[k:k+n/2]}D_{[k:k+n/2]}(X)\mathbf{F}, (d + \bar{d})n)$

7: $\mathbf{c}[k + n/2 : k + n] \leftarrow \text{MATRIXVECTMULT}(\mathbf{F}_r, a-1, k + n/2)$

If $T(n)$ is the number of operations needed for a call to Algorithm 2 with input size $a = \log_2 n$, then we will show that

$$T(n) \leq 2T(n/2) + O(t^2(d + \bar{d})n \log n).$$

This will prove the claimed runtime.

The first recursive call only requires computing $D_{[k+n/2:k+n]}(X)\mathbf{F}$ which consists of $t+1$ multiplications of degree $(d + \bar{d})n$ polynomials; this takes $O(t(d + \bar{d})n \log^2 n)$ operations. For the second recursive call, the runtime is dominated by Step 6 is matrix vector multiplication with dimension $t+1$ where each entry is a polynomial of degree $O((d + \bar{d})n)$. Hence, this step takes $O(t^2(d + \bar{d})n \log^2 n)$ many operations, as desired. \square

We remark that the analysis here is essentially equivalent to that of Lemma 3.4; both algorithms are bottlenecked by multiplication of a ranged transition matrix. The only difference is that the transition matrices for a (d, \bar{d}) -nice recurrence have degrees scaled by a factor of $(d + \bar{d})$ as shown in Lemma 2.7. Similarly, a simple modification of Lemma 3.3 shows that the pre-processing step of computing $\mathbf{T}_{[bN/2^d : bN/2^d + N/2^{d+1}]}$ for all $0 \leq d < m, 0 \leq b < 2^d$ takes $O(t^{\omega_{\mathbb{F}}}(d + \bar{d})N \log^2 N)$ operations here. Thus, we have argued the following result:

Theorem 5.3. *For any (d, \bar{d}) -nice recurrence as in (15) satisfying (16), with $O(t^{\omega_{\mathbb{F}}}(d + \bar{d})N \log^2 N)$ pre-processing operations, any \mathbf{Ab} can be computed with $O(t^2(d + \bar{d})N \log^2 N)$ operations over \mathbb{F} .*

Corollary 5.4. *For any $(1, 0)$ -nice recurrence (11) under Definition 2.2, with $O(t^{\omega_{\mathbb{F}}}N \log^2 N)$ pre-processing operations, any \mathbf{Ab} can be computed with $O(t^2N \log^2 N)$ operations over \mathbb{F} .*

6 Recurrence Extensions

We have shown how to compute $\mathbf{A}^T \mathbf{b}$ and \mathbf{Ab} for matrices \mathbf{A} defined by (11) under Definition 2.2. In this section we show how to modify the multiplication algorithms when the recurrence uses different instantiations of \otimes .

6.1 Recurrence in a modulus

We first suppose that \mathbf{A} satisfies a (d, \bar{d}) -nice recurrence defined by (11) and Definition 2.3. We will prove

Theorem 6.1. *For any \mathbf{A} satisfying a (d, \bar{d}) -nice recurrence (11) under Definition 2.3, with $O(t^{\omega_{\mathbb{F}}}(d + \bar{d})N \log^2 N)$ pre-processing operations, the products $\mathbf{A}^T \mathbf{b}$ and \mathbf{Ab} can be computed with $O(t^2(d + \bar{d})N \log^2 N)$ operations over \mathbb{F} .*

The idea here is that after factoring out terms to clear denominators of the rational fractions, the transition matrices $\mathbf{T}_{[\ell:r]}$ essentially reduce to the basic polynomial recurrence setting, with degrees scaled by $(d + \bar{d})$. This follows from Lemma 2.7.

A^Tb The modification to the algorithm for $\mathbf{c} = \mathbf{A}^T \mathbf{b}$ is straightforward. We showed in Section 3 equation (12) that \mathbf{c} is the last element of $\mathbf{P}_{0,N} \otimes \mathbf{F}$ where

$$\mathbf{P}_{\ell,r} = \sum_{i=\ell}^{r-1} \mathbf{b}[i] \mathbf{T}_{[\ell+t:i+t]}$$

(as usual, \mathbf{F} is the $t+1$ -length vector of the initial rows). The sum can be computed using Algorithm 1; we only need to analyze how the runtime changes. The following lemma is straightforward to prove inductively using equation (13) and Lemma 2.7.

Lemma 6.2. *The vector $\mathbf{P}_{\ell,r}$ can be written in the form $D_{[\ell:r]}(X)^{-1} \mathbf{P}'_{\ell,r}$, where each element of $\mathbf{P}'_{\ell,r}$ is a polynomial of degree at most $(d + \tilde{d})(r - \ell)$.*

Thus by factoring out the denominators and keeping track of them separately, the same analysis holds; the only difference is that the degrees of polynomials involved here are scaled by a factor of $d + \tilde{d}$. By Lemma 3.4, running Algorithm 1 on these $\mathbf{P}_{\ell,r}$ has the same asymptotic running time as multiplication of t^2 polynomials of degree $(d + \tilde{d})N$, taking $O(t^2(d + \tilde{d})N \log^2 N)$ operations.

The post-processing step of computing $\mathbf{P}_{0,N} \otimes \mathbf{F}$ follows similar changes. By Definition 2.3 and the above lemma, this is $D_{[0:N]}(X)^{-1} \sum_{j=0}^t \mathbf{P}'_{0,N}[j] \mathbf{F}[j] \pmod{M(X)}$. The sum is t polynomial multiplications of degree $(d + \tilde{d})N$, which requires $O(t(d + \tilde{d})N \log N)$ operations. The polynomial $D_{[0:N]}(X)$ has degree $\tilde{d}N$ and multiplying by the inverse $\pmod{M(X)}$ can be done in $\tilde{O}(N)$ operations using the fast Euclidean Algorithm [44].

This proves the $\mathbf{A}^T \mathbf{b}$ part of Theorem 6.1.

Ab To compute this product, we will factor the matrix \mathbf{A} into matrices that we have already shown admit fast matrix-vector multiplication.

As usual \mathbf{A} is associated with polynomials $f_i(X) = \sum_{j=0}^{N-1} \mathbf{A}[i,j] X^j$. By Lemma 2.6, \mathbf{f}_i is the last element of $\mathbf{T}_{[t:i+t]} \mathbf{F}$, where this product is done $\pmod{M(X)}$. Now let the roots of $M(X)$ be $\alpha_0, \dots, \alpha_{N-1}$. Consider the $N \times N$ matrix \mathbf{Z} such that $\mathbf{Z}[i,j] = f_i(\alpha_j)$. This can be factored into $\mathbf{Z} = \mathbf{A} \mathbf{V}^T$, where \mathbf{V} is the $N \times N$ Vandermonde matrix on the α_j (Definition A.1).

On the other hand, we can factor \mathbf{Z} in a different way. Let $D(X) = \prod D_i(X)$ and let $C(X)$ be its inverse $\pmod{M(X)}$. For all i , define the polynomial $g_i(X)$ which is the last element of $\mathbf{T}_{[t:i+t]}(D(X)C(X)\mathbf{F})$, where this is computed over $\mathbb{F}[X]$. Note that these polynomials have degree at most $(d + \tilde{d})N$; furthermore, $D(X)$ can be computed in time $O(dN \log^2 dN)$ by divide-and-conquer, and $C(X)$ can be computed in $\tilde{O}(N)$ operations using the fast Euclidean Algorithm [44]. But the $g_i(X)$ exactly satisfy recurrence (15) with constraint (16), so that we can run Algorithm 2. Thus by Theorem 5.3, we can efficiently multiply by the $N \times (d + \tilde{d})N$ matrix \mathbf{A}' containing the coefficients of the $g_i(X)$. Finally, note that $g_i(\alpha_j) = f_i(\alpha_j)$ by their equivalence modulo $M(X)$. Therefore $\mathbf{Z} = \mathbf{A}' \mathbf{V}'^T$, where \mathbf{V}' is the $N \times (d + \tilde{d})N$ Vandermonde matrix on the α_j .

Thus we have the factorization $\mathbf{A} = \mathbf{A}' \mathbf{V}'^T \mathbf{V}^{-T}$. The Vandermonde and inverse Vandermonde matrices have dimensions at most $(d + \tilde{d})N$ and thus can be multiplied in $O((d + \tilde{d})N \log^2 N)$ time, so each component of this factorization admits matrix-vector multiplication in order $O(t^2(d + \tilde{d})N \log^2 N)$ operations.

We remark that this factorization of \mathbf{A} can be used to perform the multiplication $\mathbf{A}^T \mathbf{b}$ as well, with the same asymptotic runtime but a worse constant factor than directly running Algorithm 1.

This finishes the proof of Theorem 6.1.

6.2 Matrix recurrence

We now suppose that \mathbf{A} is a \mathbf{R} -matrix recurrence, that is it satisfies (11) through Definition 2.4. Assume that the recurrence is $(1,0)$ -nice, so that the elements of the transition matrices $\mathbf{T}_{[\ell:r]}$ can be represented by polynomials of degree less than N ; in Appendix B we show that this assumption can always be made by changing the ring to $\mathcal{R} = \mathbb{F}[X]/(c_{\mathbf{R}}(X))$, assuming the characteristic polynomial $c_{\mathbf{R}}(X)$ is known.

The idea is that the matrix \mathbf{A} satisfying a (\mathbf{G}, \mathbf{F}) -recurrence can essentially be factored into two matrices, one of which depends only on the recurrence \mathbf{G} , the other depending only on the initial conditions \mathbf{F} and the action of the \otimes operator.

By Lemma 2.6,

$$\begin{aligned}\mathbf{f}_i &= \sum_{j=0}^t h_{i,j}^{(0)}(X) \otimes \mathbf{f}_j \\ &= \sum_{j=0}^t \sum_{\ell=0}^{N-1} \mathbf{h}_{i,j}[\ell] X^\ell \otimes \mathbf{f}_j \\ &= \sum_{j=0}^t \mathbf{h}_{i,j} \mathbf{K}_j^T.\end{aligned}$$

In the above, we are defining $\mathbf{h}_{i,j}$ to be the coefficient (row) vector of polynomial $h_{i,j}^{(0)}(X)$:

$$h_{i,j}^{(0)}(X) = \sum_{\ell=0}^{N-1} \mathbf{h}_{i,j}[\ell] X^\ell,$$

and \mathbf{K}_j to be the matrix whose ℓ th column is $X^\ell \otimes \mathbf{f}_j$.

Thus letting \mathbf{H}_j be the matrix whose i th row is $\mathbf{h}_{i,j}$, we derive the identity

$$\mathbf{A} = \sum_{j=0}^t \mathbf{H}_j \mathbf{K}_j^T. \quad (18)$$

Note in particular that every \mathbf{H}_j satisfies a (\mathbf{G}, \cdot) -recurrence with Definition 2.2, by the second part of Lemma 2.6. Also, every $\mathbf{K}_j = \mathcal{K}(\mathbf{R}, \mathbf{f}_j)$ is a Krylov matrix by Definition 7.1. For now, we assume that multiplying any of \mathbf{K}_j and \mathbf{K}_j^T by a vector is free.

Remark 6.3. In many settings, such as for orthogonal polynomial transforms, it is possible to express the starting conditions in terms of the first one, i.e. $\mathbf{f}_i = g_i(X) \otimes \mathbf{f}_0$ for $1 \leq i \leq t$. It is straightforward to show that in this case, the factorization (18) can be written with only one term in the sum.

A^Tb To compute $\mathbf{A}^T \mathbf{b}$, it suffices to compute $\sum_{j=0}^t \mathbf{K}_j \mathbf{H}_j^T \mathbf{b}$ by equation (18). Ignoring the cost of the Krylov multiplications, it suffices to compute $\mathbf{H}_j^T \mathbf{b}$ for $0 \leq j \leq t$. Using Lemma 2.6 again, $\mathbf{H}_j^T \mathbf{b}$ is just the coefficient vector of

$$\left(\sum_{i=0}^{N-1} \mathbf{b}[i] \mathbf{T}_{[t:i+t]} \right) [t, j],$$

(that is, the (t, j) -th element of the sum, which is a $(t+1) \times (t+1)$ matrix of polynomials). Computing this sum is a direct application of Algorithm 1.

An equivalent view of this algorithm is starting from Algorithm 1, which computes the last row of $(\sum \mathbf{b}[i] \mathbf{T}_{[t:i+t]}) \otimes \mathbf{F}$ (equation 12). The only change is in the post-processing step of $\otimes \mathbf{F}$, which will now be Krylov matrix multiplications.

Ab If we define $\mathbf{b}_j = \mathbf{K}_j \mathbf{b}$, it suffices to compute $\sum_{j=0}^t \mathbf{H}_j \mathbf{b}_j$ by equation (18). Because each \mathbf{H}_j satisfies the same recurrence, it suffices to modify the initialization step of Algorithm 2 (note how this is dual to $\mathbf{A}^T \mathbf{b}$, where we only need to modify the post-processing step). We only need to replace the numerator of equation (17), which becomes

$$\mathbf{F}[t-i] = \frac{\sum_{j=0}^t h_{i,j}^{(0)}(X) \mathbf{b}_j^R(X)}{D_{[t:N+t]}(X)}.$$

By the second part of Lemma 2.6, $h_{i,j}^{(0)}(X) = \delta_{i,j}$ so this initialization has no additional cost if the $\mathbf{b}_j^R(X)$ are known.

Thus both **Ab** and **A^Tb** for a recurrence under Definition 2.4 reduce directly to the **A^Tb** and **Ab** algorithm for the same recurrence under Definition 2.2 or 2.3. It remains to analyze the cost of the Krylov multiplications. Both

$\mathbf{A}^T \mathbf{b}$ and $\mathbf{A} \mathbf{b}$ require at most t multiplications by Krylov matrices (\mathbf{K}_j for the former, \mathbf{K}_j^T for the latter). Suppose that \mathbf{R} is (α, β) -Krylov efficient. Since the \mathbf{K}_j are all Krylov matrices on \mathbf{R} , the pre-processing step only needs to be performed once, followed by the multiplications, incurring a cost of $\tilde{O}(\alpha N)$ and $\tilde{O}(t\beta N)$ respectively. This discussion proves

Theorem 6.4. *Let \mathbf{A} satisfy any (d, \bar{d}) -nice recurrence (11) under Definition 2.4, where \mathbf{R} is (α, β) -Krylov efficient with known characteristic polynomial. With $\tilde{O}((t^{\omega_F}(d + \bar{d}) + \alpha)N)$ pre-processing operations, the products $\mathbf{A}^T \mathbf{b}$ and $\mathbf{A} \mathbf{b}$ can be computed with $O((t^2(d + \bar{d}) + t\beta)N)$ operations over \mathbb{F} .*

6.3 Recurrences with Error

We now modify our approach to handle a low-rank error term in our recurrence. In particular, consider the following recurrence relations:

$$\mathbf{f}_{i+1} = \sum_{j=0}^t g_{i,j}(X) \otimes \mathbf{f}_{i-j} + \sum_{\ell=1}^r c_{i,\ell}(X) \otimes \mathbf{d}_\ell \quad (19)$$

These additional vectors $\mathbf{d}_\ell \in \mathbb{F}^N$ represent a rank- r error in our recurrence. Such a recurrence is said to have *recurrence error width* of (t, r) . We say that this recurrence is (d, \bar{d}) -nice if the $\{g_{i,j}(X)\}$ are nice, and $c_{i,\ell}(X) = n_{i,\ell}(X)/D_i(X)$ where $D_i(X)$ is as in Definition 2.5 and $\deg n_{i,j} \leq d + \bar{d}$. For the remainder of this section we assume \otimes is defined by 2.2 and $(d, \bar{d}) = (1, 0)$.

The main idea here is that the rank term can be re-interpreted to increase the recurrence width instead, so that a (t, r) -error width recurrence is reduced to a modification of a standard $(t + r)$ -width recurrence. We present two approaches for solving a recurrence of this type. The first approach is a direct reduction to our standard recurrence (11), where we fold the error terms into the recurrence using ‘dummy’ polynomials. The second approach is a direct modification of Lemma 2.6, which replaces the work-horse lemma in Algorithms 1 and 2.

We can insert dummy polynomials into the sequence to capture the errors. In particular, we insert r rows consisting of the \mathbf{d}_ℓ between every t of the original \mathbf{f}_i . Now note that every row of this resulting recurrence can be expressed in terms of the previous $t + r$ rows. Thus this results in a sequence of length $O(N \frac{t+r}{t})$ that follows a recurrence of width $t + r$. We can then apply our previous algorithms to compute $\mathbf{A}^T \mathbf{b}$ and $\mathbf{A} \mathbf{b}$. Note that the problem size depends on the width N of the matrix and not the height $N(1 + r/t)$, because we avoid any computations involving rows corresponding to the dummy polynomials since we already know them. Therefore the algorithms take $O((r + t)^2 N \log^2 N)$ operations and $O((r + t)^{\omega_F} N \log^2 N)$ preprocessing. We also remark that the r error terms can be separated and each handled with a recurrence of width $t + 1$, leading to runtime and preprocessing $O(r t^2 N \log^2 N)$ and $O(t^{\omega_F} N \log^2 N)$ instead.

This results in a simple and direct reduction to our previous algorithms to handle recurrences with error, and also showcases that our basic notion of recurrences (11) is powerful enough to capture more a more complicated recurrence (19). However, this reduction is somewhat loose (i.e. a recurrence of width $t + r$ is strictly more powerful than a recurrence of width (t, r)) and we can do better.

Next, we show how to handle the error terms more precisely to achieve a better algorithm runtime. This approach will also utilize on the previous algorithms, but as a subroutine of a larger divide-and-conquer. First we will modify the work-horse lemma to handle all the error terms.

Consider a recurrence defined by (19) again. We will switch to using a presentation where we consider dummy starting conditions $\mathbf{f}_{-t-1} = \dots = \mathbf{f}_{-1} = \mathbf{0}$, and $\mathbf{f}_0, \dots, \mathbf{f}_t$ are just elements defined by the recurrence with error: e.g. $\mathbf{f}_t = \sum_{j=0}^t 0 \otimes \mathbf{f}_{t-1-j} + \mathbf{f}_t$ (compare to (19)).

In the style of the work-horse lemma, the recurrence (19) can be written as

$$\begin{bmatrix} \mathbf{f}_{i+1} \\ \vdots \\ \mathbf{f}_{i-t+1} \end{bmatrix} = \mathbf{T}_i \otimes \begin{bmatrix} \mathbf{f}_i \\ \vdots \\ \mathbf{f}_{i-t} \end{bmatrix} + \begin{bmatrix} c_{i,1}(X) & \cdots & c_{i,r}(X) \\ \vdots & \ddots & \vdots \\ 0 & \cdots & 0 \end{bmatrix} \otimes \begin{bmatrix} \mathbf{d}_1 \\ \vdots \\ \mathbf{d}_r \end{bmatrix} \quad (20)$$

for $0 \leq i < N$. Let that $c_{i,\ell}(X)$ matrix be denoted \mathbf{C}_i . Composing (20) and using $\mathbf{f}_{-t-1} = \dots = \mathbf{f}_{-1} = \mathbf{0}$, we get an

analog to the work-horse lemma

$$\begin{bmatrix} \mathbf{f}_i \\ \vdots \\ \mathbf{f}_{i-t} \end{bmatrix} = (\mathbf{T}_{[0:i]} \mathbf{C}_0 + \mathbf{T}_{[1:i]} \mathbf{C}_1 + \cdots + \mathbf{T}_{[i:i]} \mathbf{C}_i) \otimes \begin{bmatrix} \mathbf{d}_1 \\ \vdots \\ \mathbf{d}_r \end{bmatrix}$$

We can use this work-horse lemma to compute $\mathbf{A}\mathbf{b}$ and $\mathbf{A}^T \mathbf{b}$ with a divide-and-conquer algorithm. We will show $\mathbf{A}^T \mathbf{b}$ here.

Define the row vectors $\mathbf{B}_i = [\mathbf{b}[i] \quad 0 \quad \cdots \quad 0] \in \mathbb{F}^{t+1}$, so that $\mathbf{b}[i]\mathbf{f}_i = \mathbf{B}_i [\mathbf{f}_i \quad \cdots \quad \mathbf{f}_{i-t}]^T$. Therefore, the query $\mathbf{A}^T \mathbf{b}$ or $\sum \mathbf{b}[i]\mathbf{f}_i$ is equal to

$$\sum_{i=0}^{N-1} \mathbf{B}_i \sum_{j=0}^i \mathbf{T}_{[j:i]} \mathbf{C}_j \otimes \mathbf{D} = \left(\sum_{0 \leq j \leq i < N} \mathbf{B}_i \mathbf{T}_{[j:i]} \mathbf{C}_j \right) \otimes \mathbf{D}$$

where $\mathbf{D} = [\mathbf{d}_1 \quad \cdots \quad \mathbf{d}_r]^T$.

We will drop the $\otimes \mathbf{D}$ and perform it at the end, so we only care about the sum, which can be computed using divide and conquer. We can break the sum into

$$\begin{aligned} \sum_{0 \leq j \leq i < N} \mathbf{B}_i \mathbf{T}_{[j:i]} \mathbf{C}_j &= \sum_{0 \leq j \leq i < N/2} \mathbf{B}_i \mathbf{T}_{[j:i]} \mathbf{C}_j + \sum_{j < N/2, i \geq N/2} \mathbf{B}_i \mathbf{T}_{[j:i]} \mathbf{C}_j + \sum_{N/2 \leq j \leq i < N} \mathbf{B}_i \mathbf{T}_{[j:i]} \mathbf{C}_j \\ &= \sum_{0 \leq j \leq i < N/2} \mathbf{B}_i \mathbf{T}_{[j:i]} \mathbf{C}_j + \left(\sum_{i \geq N/2} \mathbf{B}_i \mathbf{T}_{[N/2:i]} \right) \left(\sum_{j < N/2} \mathbf{T}_{[j:N/2]} \mathbf{C}_j \right) + \sum_{N/2 \leq j \leq i < N} \mathbf{B}_i \mathbf{T}_{[j:i]} \mathbf{C}_j \end{aligned}$$

Note that the first sum corresponds to $\mathbf{f}_0, \dots, \mathbf{f}_{N/2-1}$ and the other two correspond to $\mathbf{f}_{N/2}, \dots, \mathbf{f}_{N-1}$. Also note that the first and last terms are equivalent to the original sum, but of half the size. Thus they can be computed with recursive calls, and we only have to worry about the middle product.

We just need to compute the left and right sides of this product. Note that the left product is exactly equation (14) which is the main problem addressed in Section 3. It can be computed in $O(t^2 N \log^2 N)$ time, after seeing the query \mathbf{b} . The right product is again equivalent to equation (14), up to transpose and different indices on the transition matrix ranges. Since the \mathbf{C}_i are known upfront, these sums can be precomputed in $O(\alpha_r N \log^2 N)$ operations, where α_r is the time it takes to multiply a $r \times t$ by $t \times t$ matrix.

Given the left and right products, their product is found by a $1 \times t$ by $t \times r$ matrix multiplication over polynomials of degree $O(N)$, which takes $O(r t N \log N)$ operations. Performing the entire divide-and-conquer algorithm incurs an additional $\log N$ factor. Thus the total runtime is $O((t^2 + r t) N \log^2 N)$.

Analogously, the product $\mathbf{A}\mathbf{b}$ can be computed by modifying Algorithm 2 with this work-horse lemma; we provide a summary of the changes. First we make a high level observation about the $\mathbf{A}^T \mathbf{b}$ algorithm. We broke the problem into three pieces, two of which were identical problems of half the size, so that it sufficed to compute the final one. We can do the same with $\mathbf{A}\mathbf{b}$. Formally, we can define \mathbf{f}'_i to be the first row of $\left(\sum_{j=0}^{N/2-1} \mathbf{T}_{[j:i+N/2]} \mathbf{C}_j \right) \otimes \mathbf{D}$, which is the part of $\mathbf{f}_{i+N/2}$ not captured by the recursive subproblem. To compute $\mathbf{A}^T \mathbf{b}$, we needed $\sum_{i=0}^{N/2-1} \mathbf{b}[i] \mathbf{f}'_i = \left(\sum_{j < N/2, i \geq N/2} \mathbf{B}_i \mathbf{T}_{[j:i]} \mathbf{C}_j \right) \otimes \mathbf{D}$, and the sum decomposed into a product where one side was exactly the problem solved in Section 3. To compute $\mathbf{A}\mathbf{b}$, we need $\langle \mathbf{f}'_i, \mathbf{b} \rangle = \langle \left(\mathbf{T}_{[N/2:N/2+i]} \sum_{j=0}^{N/2-1} \mathbf{T}_{[j:N/2]} \mathbf{C}_j \right) \otimes \mathbf{D}, \mathbf{b} \rangle$ for $0 \leq i < N/2$. Note that by identity (8), we can fold the sum into \mathbf{D} so that we need to compute $\langle \mathbf{T}_{[N/2:N/2+i]} \otimes \mathbf{D}', \mathbf{b} \rangle$ where $\mathbf{D}' = \sum_{j=0}^{N/2-1} \mathbf{T}_{[j:N/2]} \mathbf{C}_j \otimes \mathbf{D}$. This can be computed by an application of Algorithm 2 with starting polynomials \mathbf{D}' .

Therefore we have shown

Theorem 6.5. *Let α_r denote the time it takes to multiply a $r \times t$ by $t \times t$ matrix. For any $(1,0)$ -nice recurrence with error width (t, r) , with $O(t^{\alpha_r} N \log^2 N)$ pre-processing operations, any $\mathbf{A}\mathbf{b}$ or $\mathbf{A}^T \mathbf{b}$ can be computed with $O(t(r + t) N \log^2 N)$ operations over \mathbb{F} .*

Note that t_r^α is bounded by $O((1 + k/t) t^\omega)$ operations, thus the pre-processing computation can be bounded by $O(r t^{\omega-1} + t^\omega)$ operations.

We remark that this algorithm recovers the bounds in Theorems 5.3 and 3.5 when $r = O(t)$. In particular, we lose nothing by formulating a recurrence (2) as a recurrence with $t + 1$ errors: each starting condition $\mathbf{f}_j : 0 \leq j \leq t$ is an error polynomial with corresponding coefficients $c_{i,j} = \delta_{i,j}$.

6.4 Combining Extensions

It is fairly straightforward to handle a matrix satisfying a recurrence with multiple of these extensions; the techniques in this section can be applied independently. For concreteness, we present the most general recurrence we can solve:

$$\mathbf{f}_{i+1} = \left(\sum_{j=0}^t g_{i,j}(\mathbf{R}) \mathbf{f}_{i-j} \right) + \sum_{\ell=1}^r c_{i,\ell}(\mathbf{R}) \mathbf{d}_\ell \quad (21)$$

where this is (d, \bar{d}) -nice, and \mathbf{R} is (α, β) -Krylov efficient.

In Section B we show that the $g_{i,j}$ can be interpreted as being in $\mathbb{F}[X]/(c_{\mathbf{R}}(X))$ and thus this recurrence satisfies equation (19) exactly under Definition 2.4. By Section 6.3, the error terms only modify the Work-horse Lemma 2.6, with $t + r$ computations of \otimes . By Section 6.2, these computations (i.e. the effect of \mathbf{R}) can be isolated with $\tilde{O}(\alpha N)$ preprocessing and $\tilde{O}((t + r)\beta N)$ computations. We are left with a modular recurrence under Definition 2.3, which is solved in Section 6.1.

Using all these techniques gives the result

Theorem 6.6. *If \mathbf{A} is a matrix satisfying a (d, \bar{d}) -nice \mathbf{R} -matrix recurrence (with known characteristic polynomial) with error width (t, r) , and \mathbf{R} is (α, β) -Krylov efficient, then with $\tilde{O}(((d + \bar{d})(t + r)t^{\omega_{\mathbb{F}}-1} + \alpha)N)$ pre-processing, the products $\mathbf{A}^T \mathbf{b}$ and $\mathbf{A} \mathbf{b}$ for any vector \mathbf{b} can be computed with $\tilde{O}((d + \bar{d})(t + r)t + (t + r)\beta)N$ operations.*

We remark that in these reductions, the \otimes operator was first isolated and the specific properties of each instantiation is handled somewhat separately from the core algorithms. In particular, the factorization in Section 6.2 can be performed as long as the $h_{i,j}^{(k)}(X)$ defined in Lemma 2.6 are polynomials. This is always true when the ring \mathcal{R} in Definition 2.1 is $\mathbb{F}[X]$ or a quotient ring, which holds for all of our specific instantiations in Definitions 2.2-2.4. We hypothesize that there are potentially other instantiations of \otimes with useful applications that can also be solved in a similar way using our generic framework.

7 Krylov Efficiency

In Section 6.2, we show how to factor matrix recurrences into the product of a polynomial/modular recurrence and a Krylov matrix, thus reducing the runtime of a \mathbf{R} -matrix recurrence to the Krylov efficiency of \mathbf{R} .

Definition 7.1. Given a matrix $\mathbf{M} \in \mathbb{F}^{N \times N}$ and a vector $\mathbf{y} \in \mathbb{F}^N$, the *Krylov matrix of \mathbf{M} generated by \mathbf{y}* (denoted by $\mathcal{K}(\mathbf{M}, \mathbf{y})$) is the $N \times N$ matrix whose i th column for $0 \leq i < N$ is $\mathbf{M}^i \cdot \mathbf{y}$. We say that \mathbf{M} is (α, β) -Krylov efficient if for every $\mathbf{y} \in \mathbb{F}^N$, we have that $\mathbf{K} = \mathcal{K}(\mathbf{M}, \mathbf{y})$ admits the operations $\mathbf{K} \mathbf{x}$ and $\mathbf{K}^T \mathbf{x}$ (for any $\mathbf{x} \in \mathbb{F}^N$) with $\tilde{O}(\beta N)$ many operations (with $\tilde{O}(\alpha N)$ pre-processing operations).

First, we show how Krylov efficiency can be reduced to Jordan efficiency. In section 9.5, we define a matrix \mathbf{M} to have an efficient Jordan decomposition if it has a decomposition $\mathbf{M} = \mathbf{A} \mathbf{J} \mathbf{A}^{-1}$ such that \mathbf{A} and \mathbf{A}^{-1} admit super-fast matrix-vector multiplication. Suppose that \mathbf{R} is Jordan efficient. Observe that the Krylov multiplication can be expressed as

$$\begin{aligned} \mathcal{K}(\mathbf{R}, \mathbf{y}) \mathbf{x} &= \sum_{i=0}^{N-1} \mathbf{x}[i] (\mathbf{R}^i \mathbf{y}) \\ &= \left(\sum_{i=0}^{N-1} \mathbf{x}[i] \mathbf{R}^i \right) \mathbf{y} \\ &= \mathbf{x}(\mathbf{R}) \mathbf{y} \end{aligned}$$

where we define the function $\mathbf{x}(X) = \sum \mathbf{x}[i]X^i$.

Note that $\mathbf{x}(X)$ is analytic and the multi-point Hermite-type evaluation problem on it is computable in $\tilde{O}(N)$ time [38, 39]. Thus the Krylov multiplication can be performed using Lemma 9.12, and Krylov efficiency of \mathbf{R} reduces exactly to Jordan efficiency with the same complexity bounds. Therefore all Jordan-efficient matrices are also Krylov-efficient.

However, this reduction is clearly one way– Jordan efficiency is stronger than Krylov efficiency, but the latter problem has more structure that we can take advantage of. In this section, we will show that the class of banded triangular matrices are Krylov efficient by showing that the Krylov matrix itself satisfies a recurrence to which we can apply our techniques.

7.1 Krylov Efficiency of triangular banded matrices

Let \mathbf{M} be an upper triangular Δ -banded matrix, i.e. all values other than $\mathbf{M}[i, \ell]$ for $i \leq \ell < i + \Delta$ are zero. Let \mathbf{y} be an arbitrary vector and let \mathbf{K} denote the Krylov matrix of \mathbf{M} with respect to \mathbf{y} .

We will show that \mathbf{K} satisfies a rational recurrence with recurrence error width of $(\Delta, 1)$. The same results also hold for lower triangular matrices.

Define polynomials

$$f_i(X) = \sum_{j=0}^{N-1} \mathbf{K}[i, j] \cdot X^j$$

Let $\mathbf{F} = \begin{bmatrix} f_0(X) \\ \vdots \\ f_{N-1}(X) \end{bmatrix}$. We can alternatively express this as

$$\mathbf{F} = \sum_{j=0}^{N-1} \mathbf{K}[:, j] X^j = \sum_{j=0}^{N-1} (\mathbf{M}^j \mathbf{y}) X^j = \left(\sum_{j=0}^{N-1} (\mathbf{M}X)^j \right) \mathbf{y}$$

Multiplying by $\mathbf{I} - \mathbf{M}X$, we get the equation

$$(\mathbf{I} - \mathbf{M}X)\mathbf{F} = \mathbf{y} - (\mathbf{M}X)^N \mathbf{y} \quad (22)$$

Therefore it is true that

$$(\mathbf{I} - \mathbf{M}X)\mathbf{F} \equiv \mathbf{y} \pmod{X^N} \quad (23)$$

and furthermore, \mathbf{F} can be defined as the unique solution of equation (23) because $\mathbf{I} - \mathbf{M}X$ is invertible in $\mathbb{F}[X]/(X^N)$ (since it is triangular and its diagonal is comprised of invertible elements $(1 - \mathbf{M}[i, i]X)$).

But equation (23) can be interpreted as a rational recurrence with error. Explicitly expanding (23), the $f_i(X)$ satisfy

$$(1 - \mathbf{M}[i+1, i+1]X) f_{i+1}(X) \equiv \sum_{j=0}^{\Delta} \mathbf{M}[i+1, i-j]X f_{i-j}(X) + \mathbf{y}[i] \pmod{X^N}$$

Therefore the multiplications $\mathbf{K}\mathbf{x}$ and $\mathbf{K}^T \mathbf{x}$ can be computed using Theorem 6.1.

Theorem 6.1 then implies that

Theorem 7.2. *Any triangular Δ -banded matrix is $(\Delta^\omega, \Delta^2)$ -Krylov Efficient.*

8 Hierarchy of Recurrence

In this section we show that a $t+1$ -term recurrence cannot be recovered by a t -term recurrence, showing a clear hierarchy among the matrices that satisfy our recurrence. Fix an arbitrary N . We will be looking at the simplest form of our recurrence: a polynomial family f_0, \dots, f_{N-1} such that

$$f_{i+1}(X) = \sum_{j=0}^t g_{i,j}(X) f_{i-j}(X)$$

where $\deg(g_{i,j}) \leq j$. Define $P(t)$ to be all families of N polynomials that satisfy our recurrence of size t . For simplicity, we assume that all polynomials have integer coefficients. For any polynomial family $f \in P(t)$, we define the matrix $\mathbf{M}(f)$ that contains the coefficients of the polynomials as its elements, as in Section 5. To show the hierarchy of matrices, we will show that no family in $P(t)$ can approximate the mapping specified by a particular family in $P(t+1)$.

Theorem 8.1. *For every $t \geq 1$, there exists an $f \in P(t+1)$ such that for every $g \in P(t)$*

$$\|\mathbf{M}(f) \cdot \mathbf{1} - \mathbf{M}(g) \cdot \mathbf{1}\|_2^2 = \Omega\left(\frac{N}{t}\right).$$

Proof. We are going to choose f such that $f_i(X) = X^i$ if $i = k(t+1)$ for some $k \geq 0$ and $f_i(X) = 0$ otherwise. Note that $f \in P(t+1)$. Let $\mathbf{c} = \mathbf{M}(f) \cdot \mathbf{1}$. In particular,

$$\mathbf{c}[i] = \begin{cases} 1 & \text{if } i = k(t+1) \text{ for some } k \geq 0 \\ 0 & \text{otherwise} \end{cases}.$$

Fix an arbitrary $g \in P(t)$. Let $\mathbf{c}' = \mathbf{M}(g) \cdot \mathbf{1}$; note that $\mathbf{c}'[i] = g_i(1)$. Note that if any t consecutive $g_i(1), \dots, g_{i+t-1}(1)$ are 0, the t -term recurrence implies that all subsequent polynomials in family uniformly evaluate to 0 at 1. So we have two cases: (1) $g_i(1) = 0$ for all $i > N/2$ or (2) for each k , there exists an i such that $k(t+1) < i < (k+1)(t+1) \leq N/2$ and $g_i(1) \neq 0$. In the first case $\mathbf{c}'[i] = 0$ for all $i > N/2$, and $\|\mathbf{c} - \mathbf{c}'\|_2^2 \geq \frac{N}{2(t+1)}$. Similarly, in the second case $\mathbf{c}'[i] \neq 0$ and $\mathbf{c}[i] = 0$ for each of the specified i , implying once again that $\|\mathbf{c} - \mathbf{c}'\|_2^2 \geq \frac{N}{2(t+1)}$. \square

Corollary 8.2. *For every $t \geq 1$, there exists an $f \in P(t+1)$ such that for every $g \in P(t)$, $\text{rank}(\mathbf{M}(f) - \mathbf{M}(g)) = \Omega(\frac{N}{t})$.*

Proof. Let $\mathbf{H} = \mathbf{M}(f) - \mathbf{M}(g)$. Once again, we define f such that $f_i(X) = X^i$ if $i = k(t+1)$ for some $k \geq 0$ and $f_i(X) = 0$ otherwise. Once again, we have two cases. First, $\deg(g_i(X)) < i$ for all $i = k(t+1), i > N/2$. Then $\mathbf{H}[i, i] = 1$ for all $i = k(t+1), i > N/2$, implying $\text{rank}(\mathbf{H}) \geq \frac{N}{2t}$. In the second case, we rely on the degree bound of the transition polynomials; in particular, if $g_{i+1}(X) = \sum_{j=0}^t h_{i,j}(X)g_{i-j}(X)$, we bound $\deg(h_{i,j}(X)) \leq j$. If $\deg(g_i(X)) = i$ for some $i = k(t+1), i > N/2$, we know that for each value of k such that $(k+1)(t+1) < N/2$, there exists a j such that $k(t+1) < j < (k+1)(t+1)$ and $\deg(g_j(X)) = j$. For each of these j , $\mathbf{H}[j, j] \neq 0$, implying $\text{rank}(\mathbf{H}) \geq \frac{N}{2t}$. \square

9 Special Cases

9.1 Displacement Rank

The *displacement rank* of a matrix A with respect to matrices \mathbf{L}, \mathbf{R} is defined as the rank of the *error matrix*

$$\mathbf{E} = \mathbf{L}\mathbf{A} - \mathbf{A}\mathbf{R}.$$

The concept of displacement rank has been used to generalize and unify common structured matrices such as Hankel, Toeplitz, Vandermonde, and Cauchy matrices; these matrices all have low displacement ranks with respect to diagonal or shift matrices being \mathbf{L} and \mathbf{R} . Olshevsky and Shokrollahi [38] defined the confluent Cauchy-like matrices to be the class of matrices with low displacement rank with respect to Jordan form matrices; this class of matrices generalized and unified the previously mentioned common structured matrices. Our class of structured matrices captures the class of matrices with low displacement rank with respect to a more general form for \mathbf{L} and \mathbf{R} .

9.1.1 Basic conditions on \mathbf{L} and \mathbf{R}

Let \mathbf{L} be Krylov efficient and suppose that we know its characteristic polynomial $\alpha_{\mathbf{L}}(X)$. Let \mathbf{R} be triangular (throughout this section, we will assume it is upper triangular) and Δ -banded and suppose that its eigenvalues are disjoint

from \mathbf{L} 's. We will show that the columns of \mathbf{A} satisfies a standard \mathbf{L} -matrix recurrence. Suppose $\text{rank}(\mathbf{E}) = r$; we can then express any column of \mathbf{E} in terms of a basis $\mathbf{d}_0, \dots, \mathbf{d}_{r-1}$. Let $\mathbf{f}_i = \mathbf{A}[:, i]$.

$$\begin{aligned}\mathbf{L}\mathbf{f}_i - \sum_{j=i-\Delta}^i \mathbf{f}_j \mathbf{R}[j, i] &= \sum_{\ell=0}^{r-1} c_{i,\ell} \mathbf{d}_\ell \\ (\mathbf{L} - \mathbf{R}[i, i] \mathbf{I}) \mathbf{f}_i &= \sum_{j=i-\Delta}^{i-1} \mathbf{f}_j \mathbf{R}[j, i] + \sum_{\ell=0}^{r-1} c_{i,\ell} \mathbf{d}_\ell \\ \mathbf{f}_i &= \sum_{j=i-\Delta}^{i-1} (\mathbf{L} - \mathbf{R}[i, i] \mathbf{I})^{-1} \mathbf{R}[j, i] \mathbf{f}_j + \sum_{\ell=0}^{r-1} (\mathbf{L} - \mathbf{R}[i, i] \mathbf{I})^{-1} c_{i,\ell} \mathbf{d}_\ell \\ \mathbf{f}_i &= \sum_{j=i-\Delta}^{i-1} \frac{\mathbf{R}[j, i]}{(\mathbf{X} - \mathbf{R}[i, i])} \otimes \mathbf{f}_j + \sum_{\ell=0}^{r-1} \frac{c_{i,\ell}}{(\mathbf{X} - \mathbf{R}[i, i])} \otimes \mathbf{d}_\ell\end{aligned}$$

where \otimes is $p \otimes \mathbf{f} = p(\mathbf{L})\mathbf{f}$. Note that the disjoint eigenvalue assumption asserts that $\mathbf{L} - \mathbf{R}[i, i]$ is invertible for all i , hence the \mathbf{f}_i are well-defined and unique. The assumption of knowing $c_{\mathbf{L}}(\mathbf{X})$ is necessary since these recurrence coefficients are actually treated as elements of the quotient ring $\mathbb{F}[\mathbf{X}]/(c_{\mathbf{L}}(\mathbf{X}))$ and not $\mathbb{F}(\mathbf{X})$ (see Section 6.1).

The above discussion implies that \mathbf{A} is a $(0, 1)$ -nice L -matrix recurrence with error width (Δ, r) . Plugging these values into Theorem 6.6 yields running times for superfast matrix-vector multiplication by \mathbf{A} and \mathbf{A}^T . Furthermore, when \mathbf{L} is also triangular and Δ -banded, its characteristic polynomial can be computed in $\tilde{O}(N)$ time and it is $(\Delta^{\omega_F}, \Delta^2)$ -Krylov efficient by Theorem 7.2.

We also note that the conditions on \mathbf{L}, \mathbf{R} are symmetric. Let (a) and (b) be matrix properties that apply through transposition (i.e. if \mathbf{A} satisfies (a) then so does \mathbf{A}^T). Now suppose that $\mathbf{A}\mathbf{b}$ and $\mathbf{A}^T\mathbf{b}$ admit fast multiplication algorithms when \mathbf{L} satisfies (a) and \mathbf{R} satisfies (b). Then they are also efficient when \mathbf{R} satisfies (a) and \mathbf{L} satisfies (b). This is because $\mathbf{R}^T \mathbf{A}^T - \mathbf{A}^T \mathbf{L}^T = -\mathbf{E}^T$ is also rank r , let $\mathbf{L}' = \mathbf{R}^T$ and $\mathbf{R}' = \mathbf{L}^T$ which satisfy the appropriate properties so \mathbf{A}^T and $(\mathbf{A}^T)^T$ admit fast multiplication.

The above discussion implies

Theorem 9.1. *Suppose $\text{rank}(\mathbf{L}\mathbf{A} - \mathbf{A}\mathbf{R}) = r$ for \mathbf{L} that is (α, β) -Krylov efficient with known characteristic polynomial, and \mathbf{R} that is triangular and Δ -banded. Suppose furthermore that \mathbf{R} and \mathbf{L} have disjoint eigenvalues. Then we can compute $\mathbf{A}\mathbf{b}$ and $\mathbf{A}^T\mathbf{b}$ for any vector \mathbf{b} in $\tilde{O}((\Delta + r)(\Delta + \beta)N)$ operations with $\tilde{O}((\Delta^{\alpha_r} + \alpha)N)$ preprocessing.*

Corollary 9.2. *Suppose $\text{rank}(\mathbf{L}\mathbf{A} - \mathbf{A}\mathbf{R}) = r$ for \mathbf{L}, \mathbf{R} that are triangular Δ -banded, and have disjoint eigenvalues. Then we can compute $\mathbf{A}\mathbf{b}$ and $\mathbf{A}^T\mathbf{b}$ for any vector \mathbf{b} in $\tilde{O}(\Delta^2(\Delta + r)N)$ operations with $\tilde{O}((\Delta^{\omega_F}(t + r))N)$ preprocessing.*

This finishes the proof of Theorem 1.4.

We remark that the most general previous displacement rank results in literature have \mathbf{L} and \mathbf{R} in Jordan normal form, which are captured by Olshovsky and Shokrollahi [38]. The results in this section cover all \mathbf{L} and \mathbf{R} in Jordan normal form that have distinct eigenvalues. We note that it is not possible to have a single efficient algorithm for matrices with low displacement rank with respect to arbitrary \mathbf{L}, \mathbf{R} in Jordan normal form. In particular, every matrix has low displacement rank with respect to $\mathbf{L} = \mathbf{R} = \mathbf{I}$. In general, when \mathbf{L} and \mathbf{R} share eigenvalues, the equation $\mathbf{L}\mathbf{A} - \mathbf{A}\mathbf{R} = \mathbf{E}$ does not uniquely specify \mathbf{A} , and we hypothesize that a general algorithm will incur an extra term roughly corresponding to the complexity of fully specifying \mathbf{A} .

However, there are fundamental results that have \mathbf{L} and \mathbf{R} with shared eigenvalues; for example, Toeplitz-like matrices have low displacement rank with respect to $\mathbf{L} = \mathbf{R} = \mathbf{S}^T$. We can generalize our results to capture these as well.

9.1.2 Further classes of displacement rank

The equation $\mathbf{L}\mathbf{A} - \mathbf{A}\mathbf{R} = \mathbf{E}$ only has a unique solution if \mathbf{L} and \mathbf{R} have disjoint eigenvalues. To see this, suppose \mathbf{L} and \mathbf{R} have at least one shared eigenvalue; we will find a non-zero solution to $\mathbf{L}\mathbf{A} - \mathbf{A}\mathbf{R} = 0$. Let $\mathbf{L} = \mathbf{X}\mathbf{J}_0\mathbf{X}^{-1}$ and $\mathbf{R} = \mathbf{Y}\mathbf{J}_1\mathbf{Y}^{-1}$ be the Jordan decompositions of the two matrices. We then have $\mathbf{J}_0\mathbf{X}^{-1}\mathbf{A}\mathbf{Y} - \mathbf{X}^{-1}\mathbf{A}\mathbf{Y}\mathbf{J}_1 = 0$. Since \mathbf{X} and \mathbf{Y} are full rank, a non-zero solution \mathbf{A} exists if and only if there exists a nonzero solution $\mathbf{M} = \mathbf{M}(\mathbf{L}, \mathbf{R})$ to $\mathbf{J}_0\mathbf{M} - \mathbf{M}\mathbf{J}_1 = 0$. Since

\mathbf{L} and \mathbf{R} share an eigenvalue, there exists x and y such that $\mathbf{J}_0[x, x] = \mathbf{J}_1[y, y]$ and $\mathbf{J}_0[x, :], \mathbf{J}_1[:, y]$ are only non-zero on the diagonal. Then $\mathbf{M}[x, y]$ is completely free; the $[x, y]$ element of $\mathbf{J}_0\mathbf{M} - \mathbf{M}\mathbf{J}_1$ is $\mathbf{M}[x, y](\mathbf{J}_0[x, x] - \mathbf{J}_1[y, y]) = 0$. Therefore, $\mathbf{L}\mathbf{A} - \mathbf{A}\mathbf{R} = \mathbf{E}$ has multiple solutions if \mathbf{L} and \mathbf{R} have shared eigenvalues.

As such, to extend our results beyond disjoint eigenvalues, we need to allow for extra parameters to specify \mathbf{A} . In this section, we deal with cases where the extra parameters can be provided by the initial conditions of our recurrence. As an example, we rework our derivation from the previous section with $\mathbf{L} = \mathbf{R} = \mathbf{S}^T$, which corresponds to Toeplitz-like matrices. We have $\mathbf{S}^T \mathbf{f}_i - \mathbf{f}_{i-1} = \sum_{\ell=0}^{r-1} c_{i,\ell} \mathbf{d}_\ell$. This gives us the recurrence $\mathbf{f}_{i-1} = X \otimes \mathbf{f}_i - \sum_{\ell=0}^{r-1} c_{i,\ell} \mathbf{d}_\ell$, which means we can apply our matrix vector multiplication algorithm to \mathbf{A} and \mathbf{A}^T . This example illustrates a simple technique to expand the class of \mathbf{L}, \mathbf{R} that we capture: we can simply put a different element of the sequence on the left side of the recurrence equation. Also note that the free element of \mathbf{A} is specified by the initialization \mathbf{f}_0 .

If we define a matrix $\mathbf{M} = \mathbf{R} - X\mathbf{I}$ over $\mathbf{F}[X]^{N \times N}$, then we can express the relation among the \mathbf{f}_i with

$$\sum_{j=0}^N \mathbf{M}[j, i] \otimes f_j = \sum_{\ell=0}^{r-1} c_{i,\ell} \mathbf{d}_\ell \quad (24)$$

We would like to extract a Δ -width \mathbf{R} -recurrence (with error) from this relationship. This requires specifying some ordering of the \mathbf{f}_i such that the relation above can be interpreted as a Δ -width recurrence. Equivalently, we need to have permutations σ, τ such that equation (24) reduces to the recurrence

$$\mathbf{M}[\sigma(i), \tau(i)] \otimes \mathbf{f}_{\sigma(i)} = \sum_{j=1}^{\Delta} \mathbf{M}[\sigma(i-j), \tau(i)] \otimes \mathbf{f}_{\sigma(i-j)} + \sum_{\ell=0}^{r-1} c_{\tau(i),\ell} \mathbf{d}_\ell \quad (25)$$

Given a permutation σ , define a permutation matrix $\mathbf{P}(\sigma)$ such that $\mathbf{P}[i, \sigma(i)] = 1$ and 0 otherwise. Equation (24) will only reduce to recurrence (25) if $\mathbf{P}(\sigma)\mathbf{M}\mathbf{P}(\tau)$ is upper triangular and Δ -banded in rows $[\Delta + 1 : N - 1]$. In other words, for $i > \Delta$, $\mathbf{M}[\sigma(i), \tau(j)] \neq 0$ if and only if $i \leq j \leq i + \Delta$. Furthermore, the rational recurrence is only usable if the polynomial $\mathbf{M}[\sigma(i), \tau(i)]$ evaluated at \mathbf{L} results in an invertible matrix for $i > \Delta$. In this case, we will say \mathbf{M} is *pseudo similar* to a *pseudo Δ -banded upper triangular matrix*.

Going back to the example of $\mathbf{L} = \mathbf{R} = \mathbf{S}^T$, \mathbf{M} is a two-banded matrix with X on the diagonal and 1 on the superdiagonal. Define $\sigma(i) = N - i$ and $\tau(i) = i + 1$ (with $\tau(N - 1) = 0$). Then $\mathbf{P}(\sigma)\mathbf{M}\mathbf{P}(\tau)$ is a two-banded matrix with 1 on the diagonal and X on the super diagonal (ignoring the first column). With these permutations, recurrence (25) is equivalent to the recurrence derived earlier: $\mathbf{f}_{i-1} = X \otimes \mathbf{f}_i - \sum_{\ell=0}^{r-1} c_{i,\ell} \mathbf{d}_\ell$.

By applying our previously established bounds to this recurrence, we can derive an explicit runtime. Note that our computation will involve both the matrices \mathbf{L} and \mathbf{R} .

Theorem 9.3. *Suppose $\mathbf{M}(\mathbf{L}, \mathbf{R})$ is pseudo similar to a pseudo Δ -banded upper triangular matrix and \mathbf{R} is (α, β) -Krylov efficient. Let $r = \text{rank}(\mathbf{L}\mathbf{A} - \mathbf{A}\mathbf{R})$ for some matrix \mathbf{A} . Then with $\tilde{O}((\Delta^{\alpha_r} + \alpha)N)$ pre-processing, both $\mathbf{A}\mathbf{b}$ and $\mathbf{A}^T\mathbf{b}$ can be computed, for any \mathbf{b} , in $\tilde{O}((\Delta + r)(\Delta + \beta)N)$ operations.*

The techniques in this section recover all of the classic Displacement Rank results: fast matrix vector multiplication for Cauchy-like, Vandermonde-like, Toeplitz-like, and Hankel-like matrices. In addition, as illustrated in the example for Toeplitz-like matrices, the permutation techniques of this section allows for any \mathbf{L} and \mathbf{R} that are Jordan blocks. Note that for all of these classes $\Delta, \alpha, \beta = O(1)$, giving us an $\tilde{O}(rN)$ runtime.

Olshevsky and Shokrollahi discuss \mathbf{L} and \mathbf{R} in general Jordan normal form [38]. Unfortunately, they do not explicitly deal with multiple solutions to the displacement rank equation; they do claim an extension of their results could handle the general case. Similarly, for our recurrence-based approach, if we extend our rational recurrences to use pseudo-inverses and allow for extra parameters specifying \mathbf{A} to be added to the errors \mathbf{d} , we could be able to handle \mathbf{L}, \mathbf{R} sharing arbitrary eigenvalues. We leave this for future investigation.

9.2 Orthogonal Polynomials

Driscoll, Healy and Rockmore [17] found a superfast matrix vector multiplication algorithm for matrices whose rows are orthogonal polynomials. They take advantage of the three-term recurrence (1) that any family of orthogonal polynomials must satisfy. In addition to being more general, we view our algorithm as being simpler and more

direct. In fact, there are two ways in which the Driscoll et al. algorithm can be viewed as a direct application of our algorithm. The Driscoll et al. algorithm actually computes the projection of a vector \mathbf{b} onto the orthogonal polynomials $p_0(X), \dots, p_{N-1}(X)$ as defined by

$$\hat{\mathbf{b}}[i] = \sum_{j=0}^{N-1} b_j p_i(z_j).$$

Note that the algorithm combines the matrix multiplication with a multi-point evaluation, while our algorithm more directly operates on the polynomials themselves instead of their evaluations. Thus if \mathbf{A} is the matrix containing the orthogonal polynomials in its rows, then $\hat{\mathbf{b}} = (\mathbf{A}\mathbf{V}^T)\mathbf{b}$, where \mathbf{V} is the Vandermonde matrix with evaluation points z_0, \dots, z_{N-1} . In fact, it can be shown that the operations performed by the Driscoll et al. algorithm are actually equivalent to first computing the projections onto monomials $\mathbf{V}^T\mathbf{b}$, and then applying our algorithm to multiply by \mathbf{A} . Even more generally, the orthogonal polynomial transform matrix directly satisfies a \mathbf{R} -matrix recurrence (5) with $\mathbf{R} = \text{diag}\langle z_0, \dots, z_{N-1} \rangle, \mathbf{f}_0 = \mathbf{1}$. This is a more general method than the $\mathbf{A}\mathbf{V}^T$ factorization because it does not depend on the functions $f_i(X)$ being polynomials, and the factorization is just equation (18) in Section 6.2.

Orthogonal polynomials are also much more structured than our generalized polynomials, which we exemplify here by providing a simple algorithm for matrix-vector multiplication involving inverses. Suppose we have orthogonal polynomials $p_1(X), \dots, p_N(X)$ and a matrix \mathbf{A} such that $p_i(X) = \sum_{j=0}^{N-1} \mathbf{A}[i, j] X^j$. By definition [14], $\int p_i(X) p_j(X) d\mu(X) = \delta_{ij}$ for some measure $\mu(X)$. Define the moment matrix $\mathbf{M}[i, j] = \int \bar{X}^j X^i d\mu(X)$. Note that if the measure is supported on the real line, then $\mathbf{M}[i, j]$ is a function of $i + j$ and \mathbf{M} is Hankel. Similarly, if μ is supported on the complex circle, $\mathbf{M}[i, j]$ is a function of $i - j$ and \mathbf{M} is Toeplitz. Note however that \mathbf{M} may not have special structure in general.

Theorem 9.4. *If \mathbf{A} is a matrix of orthogonal polynomials with respect to the measure $\mu(X)$, and \mathbf{M} is the moment matrix over the measure, then*

$$\mathbf{A}\mathbf{M}\mathbf{A}^* = \mathbf{I}$$

Proof. Consider the following sequence of relations for $0 \leq i, j < N$:

$$\begin{aligned} (\mathbf{A}\mathbf{M}\mathbf{A}^*)[i, \ell] &= \sum_{j=0}^{N-1} \mathbf{A}[i, j] \sum_{k=0}^{N-1} \mathbf{M}[j, k] \mathbf{A}^*[k, \ell] \\ &= \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} \mathbf{A}[i, j] \left(\int X^j \bar{X}^k d\mu(X) \right) \mathbf{A}^*[k, \ell] \\ &= \int \left(\sum_{j=0}^{N-1} \mathbf{A}[i, j] X^j \right) \left(\sum_{k=0}^{N-1} \mathbf{A}^*[k, \ell] \bar{X}^k \right) d\mu(X) \\ &= \int p_i(X) \bar{p}_\ell(X) d\mu(X) \\ &= \delta_{i, \ell}, \end{aligned}$$

as desired. □

In the case that our measure is supported on the real line, $\mathbf{A}^* = \mathbf{A}^T$ and \mathbf{M} is Hankel. Note that classical superfast vector multiplication algorithms for Hankel matrices and their inverses have been well established [22]. This result implies that for orthogonal polynomials, an algorithm for the $\mathbf{A}^T\mathbf{b}$ implies an algorithm for $\mathbf{A}^{-1}\mathbf{b}$ and an algorithm for $\mathbf{A}\mathbf{b}$ implies an algorithm for $\mathbf{A}^{-T}\mathbf{b}$. So our algorithms presented in Sections 3 and 5 immediately imply algorithms for inverses in the case that our matrices contain orthogonal polynomials.

We observe that our more general class of polynomials do not follow this nice structure. $\mathbf{A}^{-1}\mathbf{A}^{-T}$ may not be a Hankel matrix in general; it may not even be symmetric.

9.3 Constant Transition Matrix

There is a great body of research relating to linear recursive sequences in which the recurrence has a constant transition matrix [35]. These sequences can be generalized from scalars to polynomials, which leads to a recurrence in which the $g_{i,j}(X)$ are independent of i . Some well-known families of polynomials fall under this setting, most prominently the various types of Chebyshev polynomials which all satisfy a width-1 recurrence with constant transitions [21]. We also note that with some transformations, other families such as the Bernoulli numbers fall under this category of recurrence, which we discuss in the next subsection.

In this case, \mathbf{T}_i as defined by Section 2 are all identical. We can then express $\mathbf{f}_i = \mathbf{T}^{i-t} \mathbf{f}$ where $\mathbf{f} = [\mathbf{f}_0 \ \cdots \ \mathbf{f}_t]^T$. In this case, we have

$$\mathbf{A}^T \mathbf{b} = \sum_{i=0}^t \mathbf{b}[i] \mathbf{f}_i + \sum_{i=t+1}^N \mathbf{b}[i] \mathbf{T}^{i-t} \mathbf{f}.$$

Similarly, we have

$$(\mathbf{Ab})[i] = \mathbf{b}^T \mathbf{T}^{i-t} \mathbf{f}.$$

We note that \mathbf{T} is essentially a traditional companion matrix of a linear recursive sequence [27], except that the coefficients in \mathbf{T} are polynomials instead of constants.

With this constant transition matrix, we can use the approach of Fiduccia [19] to reduce both the pre-processing and multiplication time to $\tilde{O}(Nt)$. Fiduccia investigates linear recurrences, and uses a fact about companion matrices to derive a $O(t \log t \log n)$ method to find any n th element of a linear recurrence with t terms. Specifically, given a companion matrix $\mathbf{C} \in \mathcal{R}^{n \times n}$ with characteristic polynomial $p(Y)$, the transformations of multiplication by \mathbf{C} in \mathcal{R}^n is isomorphic to multiplication by Y in $\mathcal{R}[Y]/(p(Y))$. Therefore computing \mathbf{C}^k can be reduced to computing $Y^k \pmod{p(Y)}$.

We can generalize this approach to our recurrence, the only change being that the elements of our companion matrix \mathbf{T} are polynomials - specifically, they are the $g_{i,j}(X)$ of (11) (which now don't depend on i , so we will denote them $g_j(X)$). Then we can define the characteristic polynomial $p(X, Y) = Y^{t+1} - \sum_{j=0}^t g_j(X) Y^{t-j}$ of \mathbf{T} . The bottlenecks of Algorithm 1 and Algorithm 2 lie in the pre-computations and multiplications of the ranged transition matrices (see Lemmas 3.4 and 5.2), so we focus our attention on them. In this case, we need to examine \mathbf{T}^{2^b} for $b \in [0, \log_2 N]$. As noted, the following formulations of jumping the recurrence by i rows are equivalent because of the isomorphism between $(\mathbf{T}, \mathbb{F}[X]^{t+1})$ and $(Y, \mathbb{F}[X][Y]/(p(X, Y)))$:

$$\begin{bmatrix} f_{k+i}(X) \\ \vdots \\ f_{k+i+t}(X) \end{bmatrix} = \mathbf{T}^i \begin{bmatrix} f_k(X) \\ \vdots \\ f_{k+t}(X) \end{bmatrix}$$

$$f_{k+i}(X) + \cdots + f_{k+i+t}(X) Y^t = Y^i (f_k(X) + \cdots + f_{k+t}(X) Y^t) \pmod{p(X, Y)} \quad (26)$$

For the preprocessing step, we will calculate polynomials $h_i(X, Y) : i = 2^b$ such that $\deg_Y(g_i) \leq t$ and $h_i(X, Y) = Y^i \pmod{p(X, Y)}$; these play the role of \mathbf{T}^i . Note that reducing a polynomial's Y -degree $\pmod{p(X, Y)}$ does not increase its overall degree:

Lemma 9.5. *Let $q(X, Y), r(X, Y) \in \mathbb{F}[X, Y]$ such that $q(X, Y) \equiv r(X, Y) \pmod{p(X, Y)}$, $\deg(q) \leq d$, and $\deg_Y(r) < \deg(p) = t + 1$. Then $\deg(r) \leq d$.*

Proof. Consider the following process of computing $r(X, Y)$ starting from $q(X, Y)$; at every iteration, we will find an equivalent polynomial $\pmod{p(X, Y)}$ but of lower Y -degree. If the polynomial has Y -degree at most t , we are done and it must be $r(X, Y)$. Otherwise, replace its term of highest Y -degree by substituting

$$Y^n = Y^{n-t-1} Y^{t+1} = Y^{n-t-1} \sum_{j=0}^t g_j(X) Y^{t-j},$$

which does not change its value $\pmod{p(X, Y)}$. It also cannot raise the polynomial's degree because of the degree bound $\deg(g_j(X)) \leq j + 1$. \square

In particular, $\deg(h_i(X, Y)) \leq i$. Given $h_i(X, Y)$, we can compute $h_{2i}(X, Y)$ by squaring $h_i(X, Y)$ and reducing (mod $p(X, Y)$). Note that $\deg_X(h_i) \leq i$ and $\deg_Y(h_i) \leq t$. Treating h_i as a polynomial in Y , squaring it takes $\tilde{O}(t)$ operations over its coefficients, which are elements of $\mathbb{F}[X]$ [12]. This in turn takes $\tilde{O}(i)$ operations over \mathbb{F} . Furthermore, computing the polynomial remainder of $h_{2i}(X, Y)$ by $p(X, Y)$ has the same complexity [13]. Thus each $h_i(X, Y)$ takes $\tilde{O}(ti)$ operations to compute, and calculating h_1, \dots, h_N through repeated squaring takes $\tilde{O}(tN)$ operations in total.

Now suppose we are given the pre-computations and are performing a matrix multiplication. For concreteness, we will examine the top level of the recurrence for $\mathbf{A}^T \mathbf{b}$. As noted in Lemma 3.4, the bottleneck was performing a matrix-vector multiplication by $\mathbf{T}_{[0:N/2]}$, which was done in $\tilde{O}(t^2 N)$ operations. In this constant-transition case, the multiplication can be performed via (26), where we can replace Y^i with the pre-computed $h_i(X, Y)$. This is again a multivariate polynomial multiplication and reduction, where the degree in X is $O(N)$ and the degree in Y is $O(t)$. Thus only $\tilde{O}(tN)$ operations are required to perform the ranged-transition multiplication. Replacing this cost in the runtime analysis of Lemma 3.4 implies that the whole algorithm only needs $\tilde{O}(tN)$ operations.

The same analysis applies for computing $\mathbf{A}\mathbf{b}$ by modifying Lemma 5.2. Finally, we incur the standard extra runtime factors for modular and matrix recurrences (Section 6.1, 6.2). Thus, we have shown that

Theorem 9.6. *For any (d, \bar{d}) -nice recurrence as in (11) where $g_{i,j}(X)$ are independent of i , with $\tilde{O}(t(d + \bar{d})N)$ pre-processing operations, any $\mathbf{A}^T \mathbf{b}$ and $\mathbf{A}\mathbf{b}$ can be computed with $\tilde{O}(t(d + \bar{d})N)$ operations over \mathbb{F} . For an \mathbf{R} -matrix recurrence (5) and \mathbf{R} is (α, β) -Krylov efficient, pre-processing and computation need additional $\tilde{O}(\alpha N)$ and $\tilde{O}(t\beta N)$ operations respectively.*

9.4 Bernoulli Polynomials

Bernoulli polynomials appear in various topics of mathematics including the Euler-Maclaurin formulae relating sums to integrals, and formulations of the Riemann zeta function [8]. In this section we will demonstrate how our techniques are flexible enough to calculate quantities such as the Bernoulli numbers, even though they do not ostensibly fall into the framework of bounded-width linear recurrences.

Bernoulli polynomials are traditionally defined by the recursive formula

$$B_i(X) = X^i - \sum_{k=0}^{i-1} \binom{i}{k} \frac{B_k(X)}{i-k+1}$$

with $B_0(X) = 1$ [6]. This recurrence seemingly cannot be captured by our model of recurrences, since each polynomial depends on *every previous polynomial*. However, with some additional work, a recurrence with bounded recurrence width can also capture this larger recurrence, thereby facilitating superfast matrix-vector multiplication involving B .

We start out by computing $B_i(0)$ for each i . For notational convenience, we may drop the 0 and use B_i to denote $B_i(0)$. In particular

$$B_i = \sum_{k=0}^i (-1)^k \frac{k!}{k+1} \left\{ \begin{matrix} i \\ k \end{matrix} \right\}$$

where $\left\{ \begin{matrix} i \\ k \end{matrix} \right\}$ denotes the Stirling numbers of the second kind [6]. To compute B_i for $0 \leq i \leq N$, we define the Stirling matrix \mathbf{W} such that $\mathbf{W}[i, j] = \left\{ \begin{matrix} i \\ j \end{matrix} \right\}$ and a vector a diagonal matrix \mathbf{D} such that $\mathbf{D}[i, i] = i!$, and a vector \mathbf{x} such that $\mathbf{x}[k] = \frac{(-1)^k k!}{k+1}$; the vector $\mathbf{b} = \mathbf{W}\mathbf{x}$ will contain $B_i(0)$ as $\mathbf{b}[i]$.

Lemma 9.7. *If $\mathbf{W}[i, j] = \left\{ \begin{matrix} i \\ j \end{matrix} \right\}$, we can multiply $\mathbf{W}\mathbf{x}$ or $\mathbf{W}^T \mathbf{x}$ for any vector \mathbf{x} with $O(N \log^2 N)$ operations.*

Proof. Note that the Stirling numbers satisfy the recurrence [6]

$$\left\{ \begin{matrix} i+1 \\ k \end{matrix} \right\} = k \left\{ \begin{matrix} i \\ k \end{matrix} \right\} + \left\{ \begin{matrix} i \\ k-1 \end{matrix} \right\}.$$

If we let $\mathbf{f}_i = \mathbf{W}[i, :]^T$, the recurrence gives us that $\mathbf{f}_{i+1} = (\mathbf{D} + \mathbf{S})\mathbf{f}_i$. Note that by Theorem 7.2, $(\mathbf{D} + \mathbf{S})$ is $(1, 1)$ -Krylov efficient. Theorems 5.3 and 3.5 complete the proof. \square

Corollary 9.8. *We can compute $B_i(0)$ for all i with $O(N \log^2 N)$ operations.*

We note that the $B_i(0)$ are actually the Bernoulli numbers, and our algorithm facilitates the computation of the Bernoulli numbers in the same runtime as recent state-of-the-art ad hoc approaches [25]. (See Section 9.6 for more.)

We now focus on a matrix \mathbf{Z} such that $\mathbf{Z}[i, j] = B_i(\alpha_j)$ for $\alpha_0, \dots, \alpha_{N-1} \in \mathbb{F}$. Note that a superfast multiplication algorithm for \mathbf{Z} immediately implies one for \mathbf{B} since $\mathbf{Z} = \mathbf{B}\mathbf{V}_\alpha^T$, where \mathbf{V}_α is the Vandermonde matrix defined by evaluation points $\alpha_0, \dots, \alpha_{N-1}$. We take advantage of the following identity relating $B_i(X)$ to B_i :

$$B_{i+1}(X) = B_{i+1} + \sum_{k=0}^i \frac{i+1}{k+1} \left\{ \begin{matrix} i \\ k \end{matrix} \right\} (X)_{k+1}$$

where $(X)_{k+1} = X(X-1)\cdots(X-k)$ denotes the falling factorial [6].

Lemma 9.9. *Suppose a matrix \mathbf{F} is defined such that $f_i(X) = \sum_{j=0}^{N-1} \mathbf{F}[i, j] X^j = (X)_i$. Then for any vector \mathbf{x} , we can multiply $\mathbf{F}\mathbf{x}$ and $\mathbf{F}^T \mathbf{x}$ in $O(N \log^2 N)$.*

Proof. By definition $f_{i+1}(X) = (X-i)f_i(X)$, and hence, Theorems 5.3 and 3.5 imply that we can multiply \mathbf{F} and \mathbf{F}^T by vectors in $O(N \log^2 N)$. \square

Theorem 9.10. *For any vector \mathbf{b} , we can compute $\mathbf{Z}\mathbf{b}$ or $\mathbf{Z}^T \mathbf{b}$ with $O(N \log^2 N)$ operations.*

Proof. Suppose we define $\mathbf{C}_{i+1, j} = \frac{i+1}{j+1} \left\{ \begin{matrix} i \\ j \end{matrix} \right\}$. Note that \mathbf{C} is just the Stirling matrix \mathbf{W} multiplied by two diagonal matrices, implying it supports $O(N \log^2 N)$ vector multiplication. Furthermore, as in our previous lemma, define \mathbf{F} to be a matrix corresponding to the falling factorials. Then $(\mathbf{C}\mathbf{V}_\alpha \mathbf{F}^T)[i, j] = B_i(\alpha_j) - B_i$. So if we define a matrix \mathbf{M} such that $\mathbf{M}[i, j] = B_i$, $\mathbf{Z} = \mathbf{C}\mathbf{V}_\alpha \mathbf{F}^T + \mathbf{M}$. Thus multiplying by \mathbf{Z} or \mathbf{Z}^T reduces to multiplying by these components, which by Lemmas 9.7 and 9.9 can be done in $O(N \log^2 N)$ operations. \square

As discussed previously, this theorem immediately implies that we can compute matrix-vector multiplications involving \mathbf{B} in $O(N \log^2 N)$ operations as well. As such, we only need to analyze the bit complexity of the algorithm to understand its numerical properties (which we do in Section 9.6).

9.5 Efficient Jordan Decomposition

Every linear operator admits a Jordan normal form, which in some senses is the smallest and most uniform representation of the operator under any basis. Knowing the Jordan decomposition of a matrix permits many useful applications. For example, being able to quickly describe and manipulate the change of basis matrix, which is a set of generalized eigenvectors, subsumes calculating and operating on the eigenvectors of the matrix.

Another use of the Jordan decomposition is being able to understand the evaluation of any analytic function on the original matrix. We highlight this application because of its connection to Krylov efficiency in Section 7, which can be solved by evaluating the matrix at a certain polynomial function. Computing matrix functions has widespread uses, and there is a significant body of research on the design and analysis of algorithms for various matrix functions, including the logarithm, matrix sign, p th root, sine and cosine [26]. Perhaps the most prominent example of a matrix function is the matrix exponential, which is tied to the theory of differential equations - systems are well-approximated around equilibria by a linearization $\mathbf{x}'(t) = \mathbf{M}\mathbf{x}(t)$, which is solved by the exponential $\mathbf{x}(t) = e^{t\mathbf{M}}\mathbf{x}(0)$ [24]. We note that much work has been done on computing the matrix exponential in particular without going through the Jordan decomposition, since it is generally hard to compute [41].

We will show how our techniques facilitate fast computation and succinct description of the Jordan decomposition for certain classes of matrices, and recapitulate how it can be used to easily compute matrix functions.

Consider a matrix \mathbf{M} with a Jordan decomposition $\mathbf{M} = \mathbf{A}\mathbf{J}\mathbf{A}^{-1}$. We say that \mathbf{M} is (α, β) -Jordan efficient if \mathbf{A} and \mathbf{A}^{-1} admit super-fast matrix-vector multiplication in $\tilde{O}(\beta N)$ operations with $\tilde{O}(\alpha N)$ pre-processing steps.

Such matrices satisfy the property that their evaluation at any analytic function f also admits super-fast matrix-vector multiplication.

Lemma 9.11 ([26]). *Let \mathbf{J} be a Jordan block with diagonal λ and $f(z)$ be a function that is analytic at λ . Then*

$$f(\mathbf{J}) = \begin{pmatrix} f(\lambda) & f'(\lambda) & \cdots & \frac{f^{(n-1)}(\lambda)}{(n-1)!} \\ 0 & f(\lambda) & \cdots & \frac{f^{(n-2)}(\lambda)}{(n-2)!} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & f(\lambda) \end{pmatrix}$$

Proof. The proof follows from considering the Taylor series expansion $f(z) = \sum \frac{f^{(i)}(\lambda)}{i!} (z - \lambda)^i$ and plugging in $\mathbf{J} = \lambda \mathbf{I} + \mathbf{S}^T$. \square

Lemma 9.12. *Let \mathbf{M} be (α, β) -Jordan efficient and $f(z)$ be a function that is analytic at the eigenvalues of \mathbf{M} . Then $f(\mathbf{M})$ admits super-fast matrix multiplication in $\tilde{O}(\beta N)$ operations, with $\tilde{O}(\alpha N)$ pre-processing.*

Proof. We can write a matrix-vector product as

$$f(\mathbf{M})\mathbf{b} = f(\mathbf{A}\mathbf{J}\mathbf{A}^{-1})\mathbf{b} = \mathbf{A}f(\mathbf{J})\mathbf{A}^{-1}\mathbf{b}$$

By Lemma 9.11, assuming access to the multi-point Hermite-type evaluations of f at the eigenvalues of \mathbf{J} , the product $f(\mathbf{J})\mathbf{b}$ is a series of Toeplitz multiplications which can be computed in time $O(N \log N)$. Thus this product is bottlenecked by the time it takes to multiply by \mathbf{A} and \mathbf{A}^{-1} , and the result follows. \square

We list some cases of Jordan-efficient matrices and show their efficiency through having low recurrence width. Note that these matrices are even more structured than general matrices of recurrence width Δ , allowing us to multiply by the inverse as well.

1. Suppose that \mathbf{M} is a Jacobi matrix. Then it is diagonalizable with $\mathbf{M} = \mathbf{A}\mathbf{D}\mathbf{A}^{-1}$, and \mathbf{A} is an orthogonal polynomial matrix. In this case, multiplication by \mathbf{A} is our standard result and multiplication by \mathbf{A}^{-1} can be done using some properties of orthogonal polynomials, described in section 9.2.
2. Suppose that \mathbf{M} is triangular Δ -banded, and \mathbf{M} 's minimal polynomial equals its characteristic polynomial. In this case, it turns out that the change of basis matrix can be expressed as $\mathbf{A}\mathbf{V}^T\mathbf{P}$, where \mathbf{A} is a Δ -width recurrence, \mathbf{V} is a confluent Vandermonde matrix, and \mathbf{P} is a permutation matrix, and furthermore $\mathbf{A}\mathbf{V}^T$ is triangular. We can show how to describe \mathbf{A} with $\tilde{O}(\Delta^\omega N)$ operations, as well as perform matrix-vector multiplication by \mathbf{A} and \mathbf{A}^{-1} in $\tilde{O}(\Delta^2 N)$ operations. The full proof is detailed below.
3. When \mathbf{M} is triangular Δ -banded and diagonal, we can also compute and multiply by \mathbf{A} efficiently using similar techniques to the above case.

Now consider a triangular Δ -banded $N \times N$ matrix \mathbf{M} whose minimal polynomial has full degree. Throughout this section, we assume \mathbf{M} is upper triangular. We will prove that \mathbf{M} is $(\Delta^\omega, \Delta^2)$ -Jordan efficient in this case.

For such a matrix \mathbf{M} , we will use $c_{\mathbf{M}}(X)$ to denote the minimal and characteristic polynomial of \mathbf{M} . Define \mathbf{F} to be the transpose of the companion matrix of $c_{\mathbf{M}}(X)$.

9.5.1 Recurrence of \mathbf{A}

\mathbf{M} is similar to $\mathbf{F}_{\mathbf{M}}$ where $\mathbf{F}_{\mathbf{M}}^T$ is the Frobenius companion matrix for $c_{\mathbf{M}}(X)$. In particular, suppose $c_{\mathbf{M}}(X) = X^N - c_{N-1}X^{N-1} - \cdots - c_0$, then $\mathbf{F}_{\mathbf{M}}$ will be

$$\begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \\ c_0 & c_1 & c_2 & \cdots & c_{N-1} \end{bmatrix}$$

This matrix $\mathbf{F}_{\mathbf{M}}^T$ is simply the Frobenius normal form, i.e., the canonical rational form, of \mathbf{M}^T .

For a matrix \mathbf{A} , we define $a_i(X)$ as the polynomial corresponding to the i^{th} row of \mathbf{A} , i.e., $a_i(X) = \sum_{j=0}^N \mathbf{A}[i, j]X^j$.

Lemma 9.13. For some matrix \mathbf{A} , $\mathbf{MA} = \mathbf{AF}_M$ if and only if $(\mathbf{M} - X\mathbf{I}) \begin{bmatrix} a_0(X) \\ \vdots \\ a_{N-1}(X) \end{bmatrix} = 0 \pmod{c_M(X)}$.

Proof. Multiplying a vector by \mathbf{F}_M in \mathcal{R} is isomorphic to multiplying the corresponding polynomial by X in $\mathcal{R}[X]/(c_M(X))$.

This implies that $\mathbf{M} \begin{bmatrix} a_0(X) \\ \vdots \\ a_{N-1}(X) \end{bmatrix} = X \begin{bmatrix} a_0(X) \\ \vdots \\ a_{N-1}(X) \end{bmatrix} \pmod{c_M(X)}$. □

Corollary 9.14. There exists a $\Delta - 1$ width matrix \mathbf{A} such that $\mathbf{MA} = \mathbf{AF}_M$. More specifically, the rows of \mathbf{A} satisfy

$$(X - \mathbf{M}[i, i])a_i(X) = \sum_{j=1}^{\min(\Delta, N-i)} \mathbf{M}[i, i+j]a_{i+j}(X) + d_i c_M(X) \quad (27)$$

where $d_i \in \mathbb{F}$.

Proof. Note that the matrix \mathbf{M} we are interested in is an upper-triangular, Δ -banded matrix. Our previous lemma

implies that that $\mathbf{MA} = \mathbf{AF}_M$ if and only if $(\mathbf{M} - X\mathbf{I}) \begin{bmatrix} a_0(X) \\ \vdots \\ a_{N-1}(X) \end{bmatrix} = 0 \pmod{c_M(X)}$.

We remove the modulus and treat the $a_i(X)$ as polynomials over $\mathbb{F}[X]$ of degree less than N . The above condition becomes the system of equations 27. Since the left side has degree N , the d_i must have degree 0 so they are scalars. Thus if a family of polynomials is defined to satisfy the recurrence above, the associated $\Delta - 1$ width matrix \mathbf{A} must satisfy $\mathbf{MA} = \mathbf{AF}_M$. □

We say that a recurrence of form (27) has *size* N if the total length of the recurrence has length N and $c(X)$ has degree N , so that all polynomials are bounded by degree $N - 1$. We call the scalars $\mathbf{M}[\cdot, \cdot]$ the recurrence coefficients, and the d_i the error coefficients.

From now on, our use of \mathbf{A} will denote such a $\Delta - 1$ width matrix. Note that independent of the d_i coefficients, the following divisibility lemma holds.

Lemma 9.15. $\prod_{j=0}^{i-1} (X - \mathbf{M}[j, j]) \mid a_i(X)$ for $0 \leq i \leq N - 1$.

Proof. Proof by induction. As a base case, this is true for $a_{N-1}(X)$ by equation (27) since $\prod_{j=0}^{N-1} (X - \mathbf{M}[j, j]) = c_M(X)$. And (27) gives us the inductive step. □

9.5.2 Conditions on the Error Coefficients

We showed in Section 9.5.1 that the equation $\mathbf{MA} = \mathbf{AF}$ is equivalent to a recurrence (27). In order to complete the task of finding \mathbf{A} satisfying $\mathbf{M} = \mathbf{AFA}^{-1}$, we must find scalars d_i in the recurrence that lead to an invertible matrix \mathbf{A} . The goal of this subsection is in proving a strong sufficiency condition on such sequences d_i .

We will first show some equivalent conditions to \mathbf{A} being invertible. We will make heavy use of (confluent) Vandermonde matrices here (Definition A.1), and for convenience define \mathbf{V}_M to be the Vandermonde matrix $\mathbf{V}_{\mathbf{M}[0,0], \dots, \mathbf{M}[N-1, N-1]}$.

Lemma 9.16. The following are true for \mathbf{A} defined by recurrence (27), for any d_i .

- (a) \mathbf{AV}_M^T is upper triangular.
- (b) $a_i^{(n_j)}(\mathbf{M}[j, j]) = 0$ for all i and $j < i$.
- (c) $\prod_{j < i} (X - \mathbf{M}[j, j]) \mid a_i(X)$ for all i .

Proof. We show the equivalence of the conditions. This is sufficient since (c) is true by Lemma 9.15.

- (a) \iff (b) As noted above, $\mathbf{A}\mathbf{V}_{\mathbf{M}}^T[i, j]$ can be understood as a Hermetian evaluations (evaluation a polynomial and its derivatives) and is equal to $a_i^{(n_j)}(\mathbf{M}[j, j])$. The equivalence follows since upper triangularity the same as saying $\mathbf{A}\mathbf{V}_{\mathbf{M}}^T[i, j] = 0$ for all $j < i$.
- (b) \iff (c) Fix an i . For every λ , let n_λ be its multiplicity in $\prod_{j < i} (X - \mathbf{M}[j, j])$. Note that condition (c) is equivalent to saying $(X - \lambda)^{n_\lambda} \mid a_i(X)$ for all λ . For a fixed λ , the divisibility condition $(X - \lambda)^{n_\lambda} \mid a_i(X)$ is equivalent to $a_i^{(k)}(\lambda) = 0$ for $k < n_\lambda$. This can be seen, for example, by considering the Taylor expansion of a_i around λ . Finally, the union of these equations over all λ is exactly (b), completing the equivalence.

□

The equivalence of the following conditions follows easily from the same reasoning.

Corollary 9.17. *The following are equivalent for \mathbf{A} defined by recurrence (27).*

- (a) \mathbf{A} is invertible.
- (b) $\mathbf{A}\mathbf{V}_{\mathbf{M}}^T$ is invertible.
- (c) $a_i^{(n_i)}(\mathbf{M}[i, i]) \neq 0$.
- (d) $\prod_{j \leq i} (X - \mathbf{M}[j, j]) \nmid a_i(X)$

We use Corollary 9.17 and in particular (d) to find conditions when \mathbf{A} is invertible. For convenience, define $p_i(X) = \prod_{j < i} \mathbf{M}[j, j]$.

Consider a fixed i . We will show that if d_{i+1}, \dots, d_{N-1} are fixed such that (d) is satisfied for all $a_{[i+1:N]}(X)$, then we can choose d_i such that (d) is satisfied for $a_i(X)$ as well.

Case 1: There does not exist $j > i$ such that $\mathbf{M}[j, j] = \mathbf{M}[i, i]$.

By the recurrence,

$$a_i(X) = \frac{\sum_{j > i} \mathbf{M}[i, j] a_j(X)}{X - \mathbf{M}[i, i]} + d_i \frac{p_N(X)}{X - \mathbf{M}[i, i]}$$

Note that the first term on the RHS is a polynomial and also is a multiple of $p_i(X)$ by inductively invoking Lemma 9.15. Note that the second term is a multiple of $p_i(X)$ but not $p_{i+1}(X)$ by the assumption for this case. Thus there exists some d_i such that $\prod_{j \leq i} (X - \mathbf{M}[j, j]) \nmid a_i(X)$ holds - in fact, there is only one d_i such that it doesn't hold.

Case 2: There exists $j > i$ such that $\mathbf{M}[j, j] = \mathbf{M}[i, i]$. Choose j to be the largest such index.

Claim 1. Fix a_j and follow the recurrence (1) without adding multiples of the modulus - i.e. let $d_{[i:j]} = 0$. Let this family of polynomials be $a'_i(X)$. Then $p_{i+1}(X) \mid a_i(X)$ if and only if $p_{i+1}(X) \mid a'_i(X)$.

Proof. $a_i(X)$ and $a'_i(X)$ differ only through the terms $c_i p_N(X), \dots, c_{j-1} p_N(X)$ added while following (1) with the modulus. All of these terms contribute a multiple of $p_{i+1}(X)$ to row i so $a_i(X) \equiv a'_i(X) \pmod{p_{i+1}(X)}$. □

Claim 2. $p_{i+1}(X) \mid a'_i(X)$ if and only if $p_{j+1}(X) \mid a_j(X)$.

Proof. The recurrence without the mod can be written as the product of $\Delta \times \Delta$ transition matrices, and in particular we can express $a'_i(X)$ via the equation

$$\frac{a'_i(X)}{p_i(X)} = h_{i,j}(X) \frac{a_j}{p_j} + h_{i,j+1}(X - \mathbf{M}[j, j]) \frac{a_{j+1}}{p_{j+1}} + \dots + h_{i,j+\Delta}(X - \mathbf{M}[j, j]) \dots (X - \mathbf{M}[j + \Delta - 1, j + \Delta - 1]) \frac{a_{j+\Delta}}{p_{j+\Delta}}$$

Multiplying through again to isolate a'_i and taking this mod $p_{i+1}(X)$, we see that only the first term affects whether $p_{i+1}(X) \mid a'_i(X)$.

Thus the claim is equivalent to saying that $h_{i,j}(X)$ is not a multiple of $(X - \mathbf{M}[i, i])$. If it was, then note that $a'_i(X)$ will *always* be a multiple of $p_{i+1}(X)$ no matter what $a_j(X)$ is. Then $a_i(X)$ will be a multiple of $p_{i+1}(X)$ no matter what a_{i+1}, \dots, a_{N-1} are, and there is *no* family of polynomials satisfying (d) of Corollary 9.17. This is a contradiction since there is some \mathbf{A} satisfying recurrence (27) that is invertible - since there is a change of basis matrix \mathbf{A} such that $\mathbf{M} = \mathbf{AFA}^{-1}$, and by Corollary 9.14 it satisfies (27). \square

The results above thus imply the following result.

Theorem 9.18. *Let \mathbf{A} be a matrix satisfying (27).*

- $\mathbf{AV}_\mathbf{M}^T$ is upper triangular for any sequences d_i .
- We can pick d_{N-1}, \dots, d_0 in order, such that each d_i chosen to satisfy the local condition $a_i^{(n_i)}(M[i, i])! = 0$. Then the matrix $\mathbf{AV}_\mathbf{M}^T$ will be invertible. Furthermore, if there exists $j > i$ such that $M[i, i] = M[j, j]$ then d_i can be anything, in particular 0.

9.5.3 Finding the Error Coefficients and Inverting $\mathbf{AV}_\mathbf{M}^T$

The main goal of this section is to prove a structure result on $\mathbf{AV}_\mathbf{M}^T$, that will easily allow us to

- Given a recurrence with unspecified error coefficients, pick the d_i such that $\mathbf{AV}_\mathbf{M}^T$ is invertible.
- Given a fully specified recurrence such that $\mathbf{AV}_\mathbf{M}^T$ is invertible, multiply $(\mathbf{AV}_\mathbf{M}^T)^{-1}\mathbf{b}$ fast.

Suppose we have fixed error coefficients $d_{[0:N]}$ and consider matrix \mathbf{A} corresponding to the polynomials generated by a recurrence of the form (27)

$$(X - \lambda_i) a_i(X) = \sum_{j=1}^{\min(\Delta, N-i)} c_{i,j} a_{i+j}(X) + d_i p_{\lambda_{[0:N]}}(X) \quad (28)$$

(We will use the shorthand notation $p_{\alpha, \beta, \dots}(X) = (X - \alpha)(X - \beta) \dots$)

By triangularity, we can partition the matrix as follows

$$\mathbf{AV}_{\lambda_{[0:N]}}^T = \begin{bmatrix} \mathbf{T}_L & \mathbf{B} \\ \mathbf{0} & \mathbf{T}_R \end{bmatrix}.$$

The core result is that the two triangular sub-blocks have the same structure as itself.

Lemma 9.19. *Given an N -size recurrence (28) that defines a matrix \mathbf{A} , there exist $N/2$ -size recurrences of the form (28) producing matrices $\mathbf{A}_L, \mathbf{A}_R$ such that*

$$\mathbf{T}_L = \mathbf{A}_L \mathbf{V}_{\lambda_{[0:N/2]}}^T \quad \text{and} \quad \mathbf{T}_R = \mathbf{A}_R \mathbf{V}_{\lambda_{[N/2:N]}}^T \mathbf{E}$$

for a matrix \mathbf{E} that is a direct sum of upper-triangular Toeplitz matrices and invertible.

Furthermore, the coefficients of these recurrences can be found with $O(\Delta^2 N \log^3 N)$ operations.

Proof. Define $c_L(X) = \prod_{i=0}^{N/2-1} (X - \lambda_i)$ and $c_R(X) = \prod_{i=N/2}^{N-1} (X - \lambda_i)$, corresponding to the left and right halves of the diagonal elements.

Structure of \mathbf{T}_R

\mathbf{T}_R consists of the higher order (derivative) evaluations of $a_{[N/2:N]}$. However, we would like to express it instead as low order evaluations (namely, corresponding to $\mathbf{V}_{\lambda_{[N/2:N]}}$) of low-degree polynomials. Fix a root λ of $p_{\lambda_{[N/2:N]}}$ and suppose that \mathbf{T}_R “contains” the evaluations $a_i^{(j)}(\lambda), \dots, a_i^{(k)}$ for $k \geq j \geq 0$. Equivalently, j is the multiplicity of λ in $\lambda_{[0:N/2]}$ and $k - j$ is the multiplicity in $\lambda_{[N/2:N]}$.

Let $q_i(X) = a_i(X)/c_L(X)$ and consider the Taylor expansions of $a_i(X)$, $q_i(X)$, $c_L(X)$ around λ

$$a_i(X) = \sum_{\ell=j}^{\infty} \frac{\alpha_i^{(\ell)}(\lambda)}{\ell!} (X-\lambda)^\ell \quad q_i(X) = \sum_{\ell=0}^{\infty} \frac{q_i^{(\ell)}(\lambda)}{\ell!} (X-\lambda)^\ell \quad c_L(X) = \sum_{\ell=j}^{\infty} \frac{c_L^{(\ell)}(\lambda)}{\ell!} (X-\lambda)^\ell$$

Since polynomial multiplication corresponds to a convolution of coefficients,

$$\begin{bmatrix} a_i^{(j)}(\lambda)/j! \\ \vdots \\ a_i^{(k)}(\lambda)/k! \end{bmatrix} = \begin{bmatrix} q_i^{(0)}(\lambda)/0! \\ \vdots \\ q_i^{(k-j)}(\lambda)/(k-j)! \end{bmatrix} * \begin{bmatrix} c_L^{(j)}(\lambda)/j! \\ \vdots \\ c_L^{(k)}(\lambda)/k! \end{bmatrix}$$

where $*$ denotes the convolution. Consider the matrix \mathbf{A}_R where row i consists of the coefficients of $q_i = a_i/c_L$. If S is the indices of the subsequence of $\lambda_{[N/2:N]}$ corresponding to λ , then the above equation can be equivalently written

$$(\mathbf{T}_R)_{i,S} = (\mathbf{A}_R \mathbf{V}_{\lambda_{[N/2:N]}}^T)_{i,S} \begin{bmatrix} c_L^{(j)}(\lambda)/j! & \cdots & c_L^{(k)}(\lambda)/k! \\ \vdots & \ddots & \vdots \\ 0 & \cdots & c_L^{(j)}(\lambda)/j! \end{bmatrix}$$

There is an analogous upper-triangular Toeplitz matrix for each λ , so define \mathbf{E} to be the $N \times N$ matrix that is the appropriate direct sum of these Toeplitz matrices, such that $(\mathbf{T}_R)_i = (\mathbf{A}_R \mathbf{V}_{\lambda_{[N/2:N]}}^T)_i \mathbf{E}$ holds. Note that the unique entries of \mathbf{E} correspond to the evaluation of c_L at the second half of $\mathbf{V}_{\lambda_{[0:N]}}$ which can be computed in $O(N \log^2 N)$. Also, \mathbf{E} is upper-triangular with non-zero diagonal. Since the above equation holds for all $i \in [N/2 : N]$, we can factor

$$\mathbf{T}_R = (\mathbf{A}_R \mathbf{V}_{\lambda_{[N/2:N]}}^T) \mathbf{E}$$

Finally, the polynomials that \mathbf{A}_R correspond to satisfy a recurrence

$$(X - \lambda_i) q_i(X) = \sum_{j=1}^{\min(\Delta, N-i)} c_{i,j} q_{i+j}(X) + d_i p_{\lambda_{[N/2:N]}}(X)$$

which directly follows from dividing (28) by $p_{\lambda_{[0:N/2]}}$.

Structure of \mathbf{T}_L

\mathbf{T}_L corresponds to evaluations of $a_{[0:N/2]}(X)$ and their derivatives at $\lambda_{[0:N/2]}$. Note that we can subtract $c_L(X)$ from any polynomial without changing these evaluations. So if we define the polynomials $q_i(X) : i \in [0 : N/2]$ to be the unique polynomials of degree less than $N/2$ such that $q_i = a_i \pmod{c_L}$ and \mathbf{A}_L to be the corresponding matrix of coefficients of q_i , then $\mathbf{T}_L = \mathbf{A}_L \mathbf{V}_{\lambda_{[0:N/2]}}^T$.

It remains to show that the q_i satisfy a recurrence of the form (28) and to specify its coefficients. Then \mathbf{A}_L will be parameterized the same way as \mathbf{A} , completing the self-similarity result that \mathbf{T}_L has the same structure as $\mathbf{A} \mathbf{V}_{\lambda_{[0:N/2]}}^T$ but of half the size.

Consider an $i \in [0 : N/2]$. Note that (28) implies that

$$\begin{aligned} (X - \lambda_i) a_i(X) &= \sum c_{i,j} a_{i+j}(X) \pmod{p_{\lambda_{[0:N]}}(X)} \\ (X - \lambda_i) a_i(X) &= \sum c_{i,j} a_{i+j}(X) \pmod{p_{\lambda_{[0:N/2]}}(X)} \\ (X - \lambda_i) q_i(X) &= \sum c_{i,j} q_{i+j}(X) \pmod{p_{\lambda_{[0:N/2]}}(X)} \end{aligned} \tag{29}$$

so that we know there exists scalars d'_i such that

$$(X - \lambda_i) q_i(X) = \sum_{j=1}^{\min(\Delta, N-i)} c_{i,j} q_{i+j}(X) + d'_i p_{\lambda_{[0:N/2]}}(X) \tag{30}$$

and the goal is to find these scalars. Define $b_i(X) = (a_i - q_i) / p_{\lambda_{[0:N/2]}}$ which is a polynomial from the definition of q_i . Subtract the previous equation from (28) and divide out $p_{\lambda_{[0:N/2]}}$ to obtain

$$(X - \lambda_i) b_i(X) = \sum_{j=1}^{\min(\Delta, N-i)} c_{i,j} b_{i+j}(X) + (d_i p_{\lambda_{[N/2:N]}}(X) - d'_i)$$

Therefore d'_i is the unique element of \mathbb{F} that makes the RHS of the above equation divisible by $(X - \lambda_i)$. This element is just $\sum c_{i,j} b_{i+j}(\lambda_i) + d_i p_{\lambda_{[N/2:N]}}(\lambda_i)$. Note that $p_{\lambda_{[N/2:N]}}$ can be evaluated at all $\lambda_{[0:N/2]}$ in time $O(N \log^2 N)$ time, so that term can be treated separately. Then define $d''_i := d'_i - d_i p_{\lambda_{[N/2:N]}}(\lambda_i)$ which is equivalently defined as the unique number making $\sum_{j=1}^{\min(\Delta, N-i)} c_{i,j} b_{i+j}(X) - d''_i$ divisible by $(X - \lambda_i)$. It suffices to find the d''_i .

Proposition 9.20. *Suppose $b_N, \dots, b_{N+\Delta}$ are polynomials of degree N and $\lambda_i, c_{i,j}$ are some fixed scalars. For $i = N-1, \dots, 0$, define b_i to be unique polynomial such that*

$$(X - \lambda_i) b_i(X) = \sum_{j=1}^{\min(\Delta, N-i)} c_{i,j} b_{i+j}(X) - d''_i$$

holds for some d''_i (which is also unique). Then we can find all the d''_i in time $O(\Delta^2 N \log^3 N)$.

Proof. If $N = O(\Delta)$, we can explicit compute the b_i and d''_i naively in time at most $O(\Delta^3)$. Otherwise we will find the d''_i with a divide-and-conquer algorithm.

As previously noted, the value of d''_i depends only on the values of $b_{[i+1:i+\Delta]}$ evaluated at λ_i . Conversely, the polynomial b_i contributes to the result only through its evaluations at $\lambda_{[i-\Delta:i-1]}$. In particular, $b_{[N/2+\Delta:N+\Delta]}$ can be reduced $(\text{mod } \prod_{[N/2:N]} (X - \lambda_j))$ without affecting the d''_i . So the remainders of the starting conditions $b_{[N:N+\Delta]} \text{ mod } \prod_{[N/2:N]} (X - \lambda_j)$ suffice to recursively compute $d''_{[N/2:N]}$. Using these, we can use the ranged transition matrix to compute $b_{[N/2:N/2+\Delta]}$. Reducing these $(\text{mod } \prod_{[0:N/2]} (X - \lambda_j))$ and recursing gives $d''_{[0:N/2]}$. The base cases need $O(\Delta^3 \cdot N / \Delta)$ operations which is dominated by the main recursion, which has runtime $T(N) = 2T(N/2) + \Delta^2 N \log^2 N$ resolving to $O(\Delta^2 N \log^3 N)$ assuming that the ranged transition matrices for the jumps are pre-computed. \square

This auxiliary algorithm gives us the following result.

Proposition 9.21. *Given the coefficients of the recurrence (28) and starting conditions $a_{[N/2:N/2+\Delta]}$, we can find the coefficients of recurrence (30) in $O(\Delta^2 N \log^3 N)$ operations.*

Proof. First find $b_{[N/2:N/2+\Delta]}$ which reduces $a_i \text{ mod } p_{\lambda_{[0:N/2]}}$. Then run the above algorithm to find all $d''_{[0:N/2]}$. Finally, set $d'_i = d''_i + d_i p_{\lambda_{[N/2:N]}}(\lambda_i)$. The bottleneck is the auxiliary algorithm which requires $O(\Delta^2 N \log^3 N)$ time. \square

Corollary 9.22. *Given the coefficients of recurrence (30), we can find coefficients of (28) in $O(\Delta^2 N \log^3 N)$ operations.*

Proof. Same as above, but the last step is reversed. Given d'_i , we compute d_i as $d_i = (d'_i - d''_i) / p_{\lambda_{[N/2:N]}}(\lambda_i)$. This is well-defined if $p_{\lambda_{[N/2:N]}}(\lambda_i)$ is non-zero. Otherwise, λ_i appears later in the sequence and we can just set $d_i = 0$ by Theorem 9.18. \square

This completes the proof of the self-similarity structure of sub-blocks of $\mathbf{A}^T_{\lambda_{[0:N/2]}}$. \square

Using Lemma 9.19, we can pick coefficients d_i for (28) that generate an invertible \mathbf{A} , and also multiply a vector by \mathbf{A}^{-1} .

Corollary 9.23. *Given a size- N recurrence (28) with fixed recurrence coefficients and starting conditions, and unspecified error coefficients d_i , we can pick the d_i such that the resulting matrix \mathbf{A} is invertible in time $O(\Delta^2 N \log^4 N)$.*

Proof. If $N = O(\Delta)$, then we can explicitly compute a_{N-1}, \dots, a_0 in order using the recurrence while picking d_{N-1}, \dots, d_0 appropriately; the results of Section 9.5.2 ensure that this is possible. Otherwise, we use Lemma 9.19 to recurse.

Since \mathbf{A}_R satisfies a size- $N/2$ recurrence with known recurrence coefficients and starting conditions and unknown error coefficients, we can recursively find $d_{[N/2:N]}$ such that \mathbf{A}_R and hence \mathbf{T}_R is invertible.

For the other half, we can jump the recurrence using $d_{[N/2:N]}$ to find $a_{[N/2:N/2+\Delta]}$ (time $O(\Delta^2 N \log^2 N)$). Now \mathbf{A}_L satisfies a size- $N/2$ recurrence with known recurrence coefficients and unknown error coefficients, and we can use [corollary inside the lemma] to find d_i such that \mathbf{A}_L is invertible.

Thus we have picked d_i such that $\mathbf{T}_L, \mathbf{T}_R$ are invertible, and by triangularity so is $\mathbf{A}\mathbf{V}_{\lambda_{[0:N/2]}}^T$.

Reducing to the subproblems is $O(\Delta^2 N \log^3 N)$ work by Lemma 9.19. \square

Corollary 9.24. *Given a fully specified recurrence (28) such that $\mathbf{A}\mathbf{V}_{\lambda_{[0:N/2]}}^T$ is invertible, we can compute $(\mathbf{A}\mathbf{V}_{\lambda_{[0:N/2]}}^T)^{-1}\mathbf{b}$ for any vector \mathbf{b} in $O(\Delta^2 N \log^4 N)$ operations.*

Proof. Note that inversion of a triangular block matrix can be expressed using the Schur complement as

$$\begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{0} & \mathbf{C} \end{bmatrix}^{-1} = \begin{bmatrix} \mathbf{A}^{-1} & -\mathbf{A}^{-1}\mathbf{B}\mathbf{C}^{-1} \\ \mathbf{0} & \mathbf{C}^{-1} \end{bmatrix}$$

So by Lemma 9.19, the desired product can be written as

$$(\mathbf{A}\mathbf{V}_{\lambda_{[0:N/2]}}^T)^{-1}\mathbf{b} = \begin{bmatrix} (\mathbf{A}_L\mathbf{V}_{\lambda_{[0:N/2]}}^T)^{-1} & -(\mathbf{A}_L\mathbf{V}_{\lambda_{[0:N/2]}}^T)^{-1}\mathbf{B}(\mathbf{A}_R\mathbf{V}_{\lambda_{[N/2:N]}}^T\mathbf{E})^{-1} \\ \mathbf{0} & (\mathbf{A}_R\mathbf{V}_{\lambda_{[N/2:N]}}^T\mathbf{E})^{-1} \end{bmatrix} \mathbf{b} = \begin{bmatrix} (\mathbf{A}_L\mathbf{V}_{\lambda_{[0:N/2]}}^T)^{-1}(\mathbf{b}_L - \mathbf{B}\mathbf{E}^{-1}(\mathbf{A}_R\mathbf{V}_{\lambda_{[N/2:N]}}^T)^{-1}\mathbf{b}_R) \\ \mathbf{E}^{-1}(\mathbf{A}_R\mathbf{V}_{\lambda_{[N/2:N]}}^T)^{-1}\mathbf{b}_R \end{bmatrix}$$

where $\mathbf{b}_L = \mathbf{b}_{[0:N/2]}$, $\mathbf{b}_R = \mathbf{b}_{[N/2:N]}$.

So $(\mathbf{A}\mathbf{V}_{\lambda_{[0:N/2]}}^T)^{-1}\mathbf{b}$ can be computed with three matrix-vector multiplications: by $(\mathbf{A}_R\mathbf{V}_{\lambda_{[N/2:N]}}^T)^{-1}$, \mathbf{E}^{-1} , \mathbf{B} , and $(\mathbf{A}_L\mathbf{V}_{\lambda_{[0:N/2]}}^T)^{-1}$, in order. Note that multiplying by \mathbf{E}^{-1} is easy since it is a direct sum of upper-triangular Toeplitz matrices. \mathbf{B} is a submatrix of \mathbf{A} which we can multiply in $\Delta^2(N \log^3 N)$ operations. Finally, the first and last multiplications are identical subproblems of half the size. \square

Thus we have shown:

Theorem 9.25. *Let \mathbf{M} be an upper-triangular, Δ -banded matrix whose minimal polynomial equals its characteristic polynomial. Then \mathbf{M} is $(\Delta^{\omega_{\mathcal{R}}}, \Delta^2)$ Jordan efficient. More precisely, there exists a Jordan decomposition $\mathbf{M} = \mathbf{A}\mathbf{J}\mathbf{A}^{-1}$ such that with $O(\Delta^{\omega_{\mathcal{R}}} N \log^2 N + \Delta^2 N \log^4 N)$ pre-processing time, multiplication by $\mathbf{A}, \mathbf{A}^{-1}$ can be computed in $O(\Delta^2 N \log^4 N)$ operations.*

9.6 Bit Complexity

It turns out that it is fairly easy to analyze the bit-complexity of our algorithms. In particular, we note that all the basic operations in our algorithms reduce to operations on polynomials. In particular, consider a matrix with recurrence width of t over the set of integers \mathbb{Z} . (We note that in most of the problems of computing sequences the input is indeed over \mathbb{Z} .) Note that our results on the efficiency of our algorithms are stated in terms of the number of operations of \mathbb{Z} . When dealing with the bit complexity of our algorithms, we have to worry about the size of the integers. It can be verified that given the input recurrence (\mathbf{G}, \mathbf{F}) , all the integers are of size $(\max(\|\mathbf{F}\|_{\infty}, \|\mathbf{G}\|_{\infty})^{O(N)})$. Since two n -bits integers can be multiplied with $O(n \log n \log \log n)$ -bit operations, this implies to obtain the bit complexity of our algorithms, we just need to multiply our bounds on the number of operations by $\tilde{O}(N \cdot \max(\|\mathbf{F}\|_{\infty}, \|\mathbf{G}\|_{\infty}))$.

In particular, the above implies that for computing the Bernoulli numbers can be computed with $\tilde{O}(N^2)$ -bit operations. This is within $\tilde{O}(1)$ factors of the best known algorithm developed specifically for this problem in [25].

9.7 Preliminary Experimental Results

We coded our basic algorithm in C++ and compared with a naive brute force algorithm. We would like to stress that in this section, we only present preliminary experimental results with the goal of showing that in principle the constants in our algorithms' runtimes are reasonable (at least to the extent that preliminary implementations of our algorithms beat the naive brute force algorithm).

We compare 3 facets of the algorithms: the preprocessing step, computing $\mathbf{A}\mathbf{b}$, and computing $\mathbf{A}^T\mathbf{b}$. We consider matrices \mathbf{A} whose rows are the coefficients of polynomials that follow our basic recurrence:

$$f_{i+1}(X) = \sum_{j=0}^t g_{i,j}(X) f_j(X)$$

where $\deg(g_{i,j}) = j$, $\deg(f_i) = i$. We generate the coefficients of $f_i(X)$ for $i \leq t$, the coefficients of $g_{i,j}(X)$ for $i > t$, and the elements of \mathbf{b} pseudo-randomly in the range $[-1, 1]$ using the C++ `rand()` function. The input to the algorithms are the $f_i(X)$ for $i \leq t$, $g_{i,j}(X)$ for $i > t$, and \mathbf{b} .

The brute force's preprocessing step is to explicitly compute the $\binom{N+1}{2}$ polynomial coefficients of $f_i(X)$ for all i , thereby computing the non-zero element of \mathbf{A} . The vector multiplication is the straightforward $O(N^2)$ algorithm. Our approach's preprocessing step uses the naive cubic matrix multiplication algorithm. And it uses the open-source library FFTW to compute FFTs.

The experiments below were run on a 2-year old laptop with an i7-4500U processor (up to 3 GHz, 4 MB cache) and 8 GB of RAM. All of the numbers below express times in seconds.

	Preprocessing		$\mathbf{A}\mathbf{b}$		$\mathbf{A}^T\mathbf{b}$	
	Brute Force	Our Approach	Brute Force	Our Approach	Brute Force	Our Approach
N = 100, t = 1	0.02	0.01	0.0004	0.003	0.0006	0.003
N = 100, t = 4	0.04	0.07	0.0005	0.009	0.0005	0.009
N = 1000, t = 1	1.3	0.13	0.05	0.04	0.06	0.04
N = 1000, t = 4	2.8	1.2	0.04	0.15	0.05	0.15
N = 10000, t = 1	127	1.7	4.8	0.4	5.6	0.5
N = 10000, t = 4	322.8	18.2	4.2	1.8	5.3	1.8

At $N = 100$, the brute force clearly outclasses ours, but notably our preprocessing isn't much slower than what brute force requires. Our approach starts to perform better at $N = 1000$ where the multiplication algorithms are similar in runtime to the brute force multiplication algorithms and are significantly faster than the brute force preprocessing. And at $N = 10000$, our approach universally outperforms the brute force approach.

10 Succinct Representations and Multivariate Polynomials

The goal of this section is two fold. The first goal is to present matrices that have low recurrence width in our sense but were not captured by previous notions of widths of structured matrices. The second goal is to show that if one can substantially improve upon the efficiency of our algorithms with respect to sharper notions of input size will leads to improvements in the state-of-the-art algorithms for multipoint evaluation of multivariate polynomials. Our initial interest in these matrices arose from their connections to coding theory, which we will also highlight as we deal with the corresponding matrices.

10.1 Multipoint evaluation of multivariate polynomials

We consider the following problem.

Definition 10.1. Given an m -variate polynomial $f(X_1, \dots, X_m)$ such that each variable has degree at most $d-1$ and $N = d^m$ distinct points $\mathbf{x}(i) = (x(i)_1, \dots, x(i)_m)$ for $1 \leq i \leq N$, output the vector $(f(\mathbf{x}(i)))_{i=1}^N$.

The best runtime for an algorithm that solves the above problem (over an arbitrary field) takes time $O(d^{\omega_2(m-1)/2+1})$, where an $n \times n$ matrix can be multiplied with an $n \times n^2$ matrix with $O(n^{\omega_2})$ operations [31,36]. (For the sake of completeness, we state these algorithms in Appendix D.) We remark on three points. First in the multipoint evaluation

problem we do not assume any structures on the N points: e.g. if the points form an m -dimensional grid, then the problem can be solved in $\tilde{O}(N)$ many operations using standard FFT techniques. Second, if we are fine with solving the problem over *finite* fields, then the breakthrough result of Kedlaya and Umans [31] solves this problem with $N^{1+o(1)}$ operations (but for arbitrary N evaluation points). In other words, the problem is not fully solved only if we do not have any structure in the evaluation points and we want our algorithms to work over arbitrary fields (or even \mathbb{R} or \mathbb{C}). Finally, from a coding theory perspective, this problem (over finite fields) corresponds to encoding of arbitrary puncturings of Reed-Muller codes.

Next, we aim to show that if we can improve our algorithms in certain settings then it would imply a fast multipoint evaluation of multivariate polynomials. In particular, we consider the following two more succinct ways of representing the input. We start with the more succinct way of representing the input. For a given polynomial $f(X) \in \mathbb{F}[X]$, let $\|f\|_0$ denote the size of the support of f . Finally, consider a matrix \mathbf{A} defined by a recurrence in (19). Define

$$\|\mathbf{A}\|_0 = \sum_{i=0}^{N-1} \sum_{j=0}^t \|g_{i,j}\|_0 + rN,$$

i.e. the size of sum of the sizes of supports of $g_{i,j}$'s plus the size of the rank r -representation of the error matrix in (19).

The second less succinct representation where we have an extra bound that $\|g_{i,j}\| \leq D$ (for potentially $D < t$) for the recurrence in (19). Then note that the corresponding matrix \mathbf{A} can be represented with size $\Theta(tDN + rN)$ elements. In this case, we will explore if one can improve upon the dependence on r in Theorem 6.6.

We would like to point out that in all of the above the way we argue that the error matrix \mathbf{E} has rank at most r is by showing it has at most r non-zero columns. Thus, for our reductions rN is also an upper bound on $\|\mathbf{E}\|_0$, so there is no hope of getting improved results in terms of the sparsity of the error matrix instead of its rank without improving upon the state-of-the-art results in multipoint evaluation of multivariate polynomials.

10.1.1 Multipoint evaluation of bivariate polynomials

We begin with the bivariate case (i.e. $m = 2$) since that is enough to connect improvements over our results to improving the state-of-the-art results in multipoint evaluation of bivariate polynomials.

For notational simplicity we assume that the polynomial is $f(X, Y) = \sum_{i=0}^{d-1} \sum_{j=0}^{d-1} f_{i,j} X^i Y^j$ and the evaluation points are $(x_1, y_1), \dots, (x_N, y_N)$. Now consider the $N \times N$ matrix

$$\mathbf{A}^{(2)} = \begin{pmatrix} 1 & x_1 & \cdots & x_1^{d-1} & y_1 & y_1 x_1 & \cdots & y_1 x_1^{d-1} & y_1^2 & y_1^2 x_1 & \cdots & y_1^2 x_1^{d-1} & \cdots & y_1^{d-1} & y_1^{d-1} x_1 & \cdots & y_1^{d-1} x_1^{d-1} \\ 1 & x_2 & \cdots & x_2^{d-1} & y_2 & y_2 x_2 & \cdots & y_2 x_2^{d-1} & y_2^2 & y_2^2 x_2 & \cdots & y_2^2 x_2^{d-1} & \cdots & y_2^{d-1} & y_2^{d-1} x_2 & \cdots & y_2^{d-1} x_2^{d-1} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_N & \cdots & x_N^{d-1} & y_N & y_N x_N & \cdots & y_N x_N^{d-1} & y_N^2 & y_N^2 x_N & \cdots & y_N^2 x_N^{d-1} & \cdots & y_N^{d-1} & y_N^{d-1} x_N & \cdots & y_N^{d-1} x_N^{d-1} \end{pmatrix}.$$

Note that to solve the multipoint evaluation problem we just need to solve $\mathbf{A}^{(2)} \cdot \mathbf{f}$, where \mathbf{f} contains the coefficients of $f(X, Y)$. Let \mathbf{D}_X and \mathbf{D}_Y denote the diagonal matrices with $\mathbf{x} = (x_1, \dots, x_N)$ and $\mathbf{y} = (y_1, \dots, y_N)$ on their diagonals respectively. Finally, define $\mathbf{Z} = \mathbf{S}^T$. Now consider the matrix

$$\mathbf{B}^{(2)} = \mathbf{D}_X^{-1} \mathbf{A}^{(2)} - \mathbf{A}^{(2)} \mathbf{Z}.$$

It can be checked that $\mathbf{B}^{(2)}$ has rank at most d . Indeed note that

$$\mathbf{B}^{(2)} = \begin{pmatrix} \frac{1}{x_1} & 0 & \cdots & 0 & \frac{y_1}{x_1} - x_1^{d-1} & 0 & \cdots & 0 & y_1 \left(\frac{y_1}{x_1} - x_1^{d-1} \right) & 0 & \cdots & 0 & \cdots & y_1^{d-2} \left(\frac{y_1}{x_1} - x_1^{d-1} \right) & 0 & \cdots & 0 \\ \frac{1}{x_2} & 0 & \cdots & 0 & \frac{y_2}{x_2} - x_2^{d-1} & 0 & \cdots & 0 & y_2 \left(\frac{y_2}{x_2} - x_2^{d-1} \right) & 0 & \cdots & 0 & \cdots & y_2^{d-2} \left(\frac{y_2}{x_2} - x_2^{d-1} \right) & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{1}{x_N} & 0 & \cdots & 0 & \frac{y_N}{x_N} - x_N^{d-1} & 0 & \cdots & 0 & y_N \left(\frac{y_N}{x_N} - x_N^{d-1} \right) & 0 & \cdots & 0 & \cdots & y_N^{d-2} \left(\frac{y_N}{x_N} - x_N^{d-1} \right) & 0 & \cdots & 0 \end{pmatrix}.$$

The above was already noticed in [37]. The above is not quite enough to argue what we want so we make the following stronger observation. Consider

$$\mathbf{C}^{(2)} = \mathbf{D}_Y^{-1} \mathbf{B}^{(2)} - \mathbf{B}^{(2)} \mathbf{Z}^d = \mathbf{D}_Y^{-1} \mathbf{D}_X^{-1} \mathbf{A}^{(2)} - \mathbf{D}_Y^{-1} \mathbf{A}^{(2)} \mathbf{Z} - \mathbf{D}_X^{-1} \mathbf{A}^{(2)} \mathbf{Z}^d - \mathbf{A}^{(2)} \mathbf{Z}^{d+1}. \quad (31)$$

One can re-write the above recurrence as follows (where $\mathbf{f}_i = (\mathbf{A}^{(2)}[i,:])^T$ and recall $\mathbf{Z} = \mathbf{S}^T$) for any $0 \leq i < N$:

$$\left(\frac{1}{x_i y_i} - \frac{\mathbf{S}}{y_i} - \frac{\mathbf{S}^d}{x_i} - \mathbf{S}^{d+1} \right) \cdot \mathbf{f}_i = (\mathbf{C}^{(2)}[i,:])^T.$$

We now claim that the rank of $\mathbf{C}^{(2)}$ is at most two. Indeed, note that

$$\mathbf{C}^{(2)} = \begin{pmatrix} \frac{1}{x_1 y_1} & 0 & \dots & 0 & \frac{-x_1^{d-1}}{y_1} & 0 & \dots & 0 & 0 & 0 & \dots & 0 & \dots & 0 & 0 & \dots & 0 \\ \frac{1}{x_2 y_2} & 0 & \dots & 0 & \frac{-x_2^{d-1}}{y_2} & 0 & \dots & 0 & 0 & 0 & \dots & 0 & \dots & 0 & 0 & \dots & 0 \\ & & & & \vdots & & & & & & & & \vdots & & & \\ \frac{1}{x_N y_N} & 0 & \dots & 0 & \frac{-x_N^{d-1}}{y_N} & 0 & \dots & 0 & 0 & 0 & \dots & 0 & \dots & 0 & 0 & \dots & 0 \end{pmatrix}.$$

Thus, we have a recurrence with recurrence width $(1, 2)$ that is $(D, 1)$ nice. Theorem 6.6 implies that we can solve the above problem with $\tilde{O}(d^3)$ operations. The algorithm of [36] uses $\tilde{O}(d^{\omega_2/2+1})$ many operations. However, note that

$$\|\mathbf{A}^{(2)}\|_0 = \Theta(d^2).$$

Thus, we have the following result:

Theorem 10.2. *If one can solve $\mathbf{A}\mathbf{b}$ for any \mathbf{b} with $\tilde{O}((\|\mathbf{A}\|_0)^{\omega_2/4+1/2-\epsilon})$ operations, then one will have an multipoint evaluation of bivariate polynomials with $\tilde{O}(d^{\omega_2/2+1-2\epsilon})$ operations, which would improve upon the currently best-known algorithm for the latter.*

10.1.2 Multipoint evaluation of multivariate polynomials

We now consider the general multivariate polynomial case. Note that we can represent the multipoint evaluation of the m -variate polynomial $f(X_1, \dots, X_m)$ as $\mathbf{A}^{(m)}\mathbf{f}$, where \mathbf{f} is the vector of coefficients and $\mathbf{A}^{(m)}$ is presented as follows.

Each of the d^m columns are indexed by tuples $\mathbf{i} \in \mathbb{Z}_d^m$ and the columns are sorted in lexicographic increasing order of the indices. Note that this implies that the $\mathbf{i} = (i_1, \dots, i_m) \in \mathbb{Z}_d^m$ column is represented by

$$\mathbf{A}^{(m)}[:, \mathbf{i}] = \begin{pmatrix} \prod_{j=1}^m x(1)_j^{i_j} \\ \prod_{j=1}^m x(2)_j^{i_j} \\ \vdots \\ \prod_{j=1}^m x(N)_j^{i_j} \end{pmatrix},$$

where the evaluation points are given by $\mathbf{x}(1), \dots, \mathbf{x}(N)$.

For notational simplicity, we will assume that m is even. (The arguments below can be easily modified for odd m .) Define recursively for $0 \leq j \leq m/2$:

$$\mathbf{B}^{(j)} = \mathbf{D}_{X_{m-j}}^{-1} \mathbf{B}^{(j+1)} - \mathbf{B}^{(j+1)} \mathbf{Z}^{d^j}, \quad (32)$$

where \mathbf{D}_{X_k} is the diagonal matrix with $(x(1)_k, \dots, x(N)_k)$ on its diagonal. Finally, for the base case we have

$$\mathbf{B}^{(\frac{m}{2}+1)} = \mathbf{A}^{(m)}.$$

It can be verified (e.g. by induction) that the recurrence in (32) can be expanded out to

$$\mathbf{B}^{(0)} = \sum_{S \subseteq [m/2, m]} (-1)^{m/2+1-|S|} \left(\prod_{j \in S} \mathbf{D}_{X_j}^{-1} \right) \mathbf{A}^{(m)} \left(\prod_{j \in [m/2, m] \setminus S} \mathbf{Z}^{d^{m-j}} \right). \quad (33)$$

The above can be re-written as (where $\mathbf{f}_i = (\mathbf{A}^{(m)})^T$):

$$\left(\sum_{S \subseteq [m/2, m]} (-1)^{m/2+1-|S|} \cdot \frac{1}{\prod_{j \in S} x(i)_j} \left(\prod_{j \in [m/2, m] \setminus S} \mathbf{s}^{d^{m-j}} \right) \right) \cdot \mathbf{f}_i = (\mathbf{B}^{(0)}[i, :])^T.$$

We will argue that

Lemma 10.3. $\mathbf{B}^{(0)}$ has rank at most $2^{m/2} \cdot d^{m/2-1}$.

Note that the above lemma implies that the recurrence in (33) is a \mathbf{S} -matrix recurrence with recurrence width $(1, r = 2^{m/2} d^{m/2-1})$ that is $(D = \frac{d^{1+m/2}-1}{d-1}, 1)$ nice. Note that in this case we have $tDN + rN = \Theta((2d)^{3m/2-1})$. Thus, we have the following result:

Theorem 10.4. *If for an \mathbf{S} -dependent recurrence we could improve the algorithm from Theorem 6.6 to run with $\tilde{O}(\text{poly}(t) \cdot DN + rN)$ operations for matrix vector multiplication, then we would be able to solve the general multipoint evaluation of multivariate polynomials in time $\tilde{O}((2d)^{3m/2-1})$, which would be a polynomial improvement over the current best algorithm (when $d = \omega(1)$), where currently we still have $\omega_2 > 3$.*

Note that the above shows that improving the dependence in r in Theorem 6.6 significantly (even to the extent of having some dependence on Dr) will improve upon the current best-known algorithms (unless $\omega_2 = 3$).

10.1.3 Proof of Lemma 10.3

We now prove Lemma 10.3. We will argue that all but at most $2^{m/2} d^{m/2-1}$ columns of $\mathbf{B}^{(0)}$ are $\mathbf{0}$, which would prove the result. Towards this end we will use the expression in (33).

For notational convenience for any index $\mathbf{i} \in \mathbb{Z}_d^m$, we will use $X^{\mathbf{i}}$ to denote the monomial $\prod_{j=1}^m X_j^{i_j}$. Then note that $\mathbf{A}^{(m)}[:, \mathbf{i}]$ is just the evaluation of the monomial $X^{\mathbf{i}}$ on the points $\mathbf{x}(1), \dots, \mathbf{x}(N)$. Further, it can be checked (e.g. by induction) that there exists polynomials $P_1(X_1, \dots, X_m)$ such that $\mathbf{B}^{(0)}[:, \mathbf{i}]$ is the evaluation of $P_1(X_1, \dots, X_m)$ on the points $\mathbf{x}(1), \dots, \mathbf{x}(N)$.

To simplify subsequent expressions, we introduce few more notation. For any $S \subseteq [m/2, m]$, define the matrix

$$\mathbf{M}_S = \left(\prod_{j \in S} \mathbf{D}_{X_j}^{-1} \right) \mathbf{A}^{(m)} \left(\prod_{j \in [m/2, m] \setminus S} \mathbf{Z}^{d^{m-j}} \right). \quad (34)$$

We can again argue by induction that for every $\mathbf{i} \in \mathbb{Z}_d^m$, we have that $\mathbf{M}_S[:, \mathbf{i}]$ is the evaluation of a polynomial $Q_S^{\mathbf{i}}(X_1, \dots, X_m)$ on the points $\mathbf{x}(1), \dots, \mathbf{x}(N)$. Note that this along with (33), implies that

$$P_1(X_1, \dots, X_m) = \sum_{S \subseteq [m/2, m]} (-1)^{m/2+1-|S|} Q_S^{\mathbf{i}}(X_1, \dots, X_m). \quad (35)$$

Now we claim that

Claim 3. For every $\mathbf{i} \in \mathbb{Z}_d^m$ that has an index $m/2 \leq j^* \leq m$ such that $i_{j^*} \geq 2$ and $S \subseteq [m/2, m] \setminus \{j^*\}$, the following holds:

$$Q_S^{\mathbf{i}}(X_1, \dots, X_m) = Q_{S \cup \{j^*\}}^{\mathbf{i}}(X_1, \dots, X_m).$$

We first argue why the claim above completes the proof. Fix any $\mathbf{i} \in \mathbb{Z}_d^m$ that has an index $m/2 \leq j^* \leq m$ such that $i_{j^*} \geq 2$. Indeed by pairing up all $S \subseteq [m/2, m] \setminus \{j^*\}$ with $S \cup \{j^*\}$, Claim 3 along with (35) implies that

$$P_1(X_1, \dots, X_m) \equiv 0.$$

Note that this implies that $\mathbf{B}^{(0)}[:, \mathbf{i}] = \mathbf{0}$ if there exists an index $m/2 \leq j^* \leq m$ such that $i_{j^*} \geq 2$. Note that there are at least $d^m - 2^{m/2} \cdot d^{m/2-1}$ such indices, which implies that at most $2^{m/2} \cdot d^{m/2-1}$ non-zero columns in $\mathbf{B}^{(0)}$, as desired.

Proof of Claim 3. For any subset $T \subseteq [m]$, recall that \mathbf{e}_T is the characteristic vector of T in $\{0, 1\}^m$. Then note that for any $S \subseteq [m/2, m]$, we have

$$Q_S^{\mathbf{i}}(X_1, \dots, X_m) = X^{(\mathbf{i} \ominus \mathbf{e}_{[m/2, m] \setminus S}) - \mathbf{e}_S},$$

where \ominus is subtraction over \mathbb{Z}_d^m (and $-$ is the usual subtraction over \mathbb{Z}^m). Indeed the $\ominus \mathbf{e}_{[m/2, m] \setminus S}$ term corresponds to the matrix $\left(\prod_{j \in [m/2, m] \setminus S} \mathbf{Z}^{d^{m-j}}\right)$ and the $-\mathbf{e}_S$ term corresponds to the matrix $\left(\prod_{j \in S} \mathbf{D}_{X_j}^{-1}\right)$ in (34).

Fix a $\mathbf{i} \in \mathbb{Z}_d^m$ that has an index $m/2 \leq j^* \leq m$ such that $i_{j^*} \geq 2$ and $S \subseteq [m/2, m] \setminus \{j^*\}$. Define $S' = S \cup \{j^*\}$. Then we claim that

$$(\mathbf{i} \ominus \mathbf{e}_{[m/2, m] \setminus S}) - \mathbf{e}_S = (\mathbf{i} \ominus \mathbf{e}_{[m/2, m] \setminus S'}) - \mathbf{e}_{S'},$$

which is enough to prove the claim. To prove the above, we will argue that

$$(\mathbf{i} \ominus \mathbf{e}_{[m/2, m] \setminus S}) = (\mathbf{i} \ominus \mathbf{e}_{[m/2, m] \setminus S'}) - \mathbf{e}_{j^*}.$$

Note that the above is sufficient by definition of S' . Now note that $\mathbf{e}_{[m/2, m] \setminus S} = \mathbf{e}_{[m/2, m] \setminus S'} + \mathbf{e}_{j^*}$. In particular, this implies that it is sufficient to prove

$$(\mathbf{i} \ominus \mathbf{e}_{[m/2, m] \setminus S'}) \ominus \mathbf{e}_{j^*} = (\mathbf{i} \ominus \mathbf{e}_{[m/2, m] \setminus S'}) - \mathbf{e}_{j^*}. \quad (36)$$

By the assumption that $i_{j^*} \geq 2$, we have that

$$(\mathbf{i} \ominus \mathbf{e}_{[m/2, m] \setminus S'})_{j^*} \geq 1,$$

which implies that both $\ominus \mathbf{e}_{j^*}$ and $-\mathbf{e}_{j^*}$ have the same effect on $(\mathbf{i} \ominus \mathbf{e}_{[m/2, m] \setminus S'})$, which proves (36), as desired. \square

10.2 Multipoint evaluation of multivariate polynomials and their derivatives

In this section, we present another example of a matrix with our general notion of recurrence that has been studied in coding theory. We would like to stress that currently this section does not yield any conditional “lower bounds” along the lines of Theorem 10.2 or 10.4.

We begin by setting up the notation for the derivative of a multivariate polynomial $f(X_1, \dots, X_m)$. In particular, given an $\iota = (i_1, \dots, i_m)$, we denote the ι th derivative of f as follows:

$$f^{(\iota)}(X_1, \dots, X_m) = \frac{\partial^{i_1}}{\partial X_1} \cdots \frac{\partial^{i_m}}{\partial X_m} f,$$

where $\frac{\partial^0}{\partial X} f = f$. If the underlying field \mathbb{F} is a finite field, then the derivatives are defined as the *Hasse* derivatives.

We consider the following problem.

Definition 10.5. Given an m -variate polynomial $f(X_1, \dots, X_m)$ such that each variable has degree at most $d-1$, an integer $0 \leq r < d$ and $n = d^m$ distinct points $\mathbf{a}(i) = (a(i)_1, \dots, a(i)_m)$ for $1 \leq i \leq N$, output the vector

$$(f^{(\iota)}(\mathbf{a}(j)))_{j \in [n], \iota \in \mathbb{Z}_r^m}.$$

The above corresponds to (puncturing) of multivariate multiplicity codes, which have been studied recently in coding theory [32–34]. These codes have excellent local and list decoding properties. We note that in definition of multiplicity codes, d and r are limits on the total degree and the total order of derivatives while in our case these are bounds for individual variables. However, these change the problem size by only a factor that just depends on m (i.e. we have at most a factor $m!$ more rows and columns). Since we think of m as constant, we ignore this difference. For such codes in [32, 34] the order of the derivatives r is assumed to be a constant. However, the multiplicity code used in [33], r is non-constant. We would like to mention that these matrices turn up in list decoding of Reed-Solomon and related codes (this was also observed in [37]).⁸ In particular, for the Reed-Solomon list decoder of Guruswami and Sudan [23], the algorithm needs \mathbf{A} with $m = 2$ and r being a polynomial in n .

⁸However in the list decoding applications \mathbf{A} is not square and one is interested in obtaining a non-zero element \mathbf{f} from its kernel: i.e. a non-zero \mathbf{f} such that $\mathbf{A}\mathbf{f} = \mathbf{0}$.

We note that the problem above is the same as $\mathbf{A} \cdot \mathbf{f}$, where \mathbf{f} is the vector of coefficients of $f(X_1, \dots, X_m) = \sum_{J \in \mathbb{Z}_d^m} f_J X^J$, where as before X^J is the monomial $\prod_{\ell=1}^m X_j^{j_\ell}$ and the matrix \mathbf{A} is defined as follows:

$$A_{(k, \iota), J} = \prod_{\ell=1}^m \binom{j_\ell}{i_\ell} (a(k)_\ell)^{j_\ell - i_\ell},$$

for $k \in [n]$, $J = (j_1, \dots, j_m) \in \mathbb{Z}_d^m$ and $\iota = (i_1, \dots, i_m) \in \mathbb{Z}_r^m$. We note that the definition holds over all fields.

We use the convention that $\binom{b}{c} = 0$ if $b < c$.

The aim of the rest of the section is to show that the matrix \mathbf{A} satisfies a recurrence that we can handle.

For notational simplicity, let us fix $k \in [n]$ and we drop the dependence on k from the indices. In particular, consider the (submatrix):

$$A_{\iota, J} = \prod_{\ell=1}^m \binom{j_\ell}{i_\ell} (a_\ell)^{j_\ell - i_\ell}.$$

Think of $\iota = (\iota^0, \iota^1)$, where $\iota^0 = (i_1, \dots, i_{m'})$ and $\iota^1 = (i_{m'+1}, \dots, i_m)$ for some $1 \leq m' < m$ to be determined. Now fix an $\iota = (\iota^0, \iota^1)$ such that $\iota^1 \neq \mathbf{0}$ and define

$$S_\iota = \{\ell \in \{m' + 1, \dots, m\} \mid i_\ell \neq 0\}.$$

Consider the following sequence of relations:

$$A_{\iota, J} = \left(\prod_{\ell \in [m] \setminus S_\iota} \binom{j_\ell}{i_\ell} (a_\ell)^{j_\ell - i_\ell} \right) \cdot \left(\prod_{\ell \in S_\iota} \left(\binom{j_\ell - 1}{i_\ell - 1} + \binom{j_\ell - 1}{i_\ell} \right) a_\ell^{j_\ell - i_\ell} \right) \quad (37)$$

$$= \sum_{T \subseteq S_\iota} \prod_{\ell=1}^m \binom{j_\ell - \mathbb{1}_{\ell \in S_\iota}}{i_\ell - \mathbb{1}_{\ell \in T}} a_\ell^{j_\ell - i_\ell} \quad (38)$$

$$= \sum_{T \subseteq S_\iota} \left(\prod_{\ell=1}^m a_\ell^{\mathbb{1}_{\ell \in S_\iota} - \mathbb{1}_{\ell \in T}} \right) \cdot \left(\prod_{\ell=1}^m \binom{j_\ell - \mathbb{1}_{\ell \in S_\iota}}{i_\ell - \mathbb{1}_{\ell \in T}} a_\ell^{j_\ell - \mathbb{1}_{\ell \in S_\iota} - (i_\ell - \mathbb{1}_{\ell \in T})} \right) \\ = \sum_{T \subseteq S_\iota} \mathbf{a}^{\mathbf{e}_{S_\iota} - \mathbf{e}_T} \cdot A_{\iota - \mathbf{e}_T, J - \mathbf{e}_{S_\iota}}. \quad (39)$$

In the above (37) follows from the following equality for integers $b \geq 0$ and $c \geq 1$, $\binom{b}{c} = \binom{b-1}{c-1} + \binom{b-1}{c}$ while (38) follows from the notation that $\mathbb{1}_P$ is the indicator value for the predicate P .

Consider the matrix \mathbf{E} defined as follows:

$$\mathbf{E}[(k, \iota), :] = \mathbf{A}[(k, \iota), :] - \sum_{T \subseteq S_\iota} \mathbf{a}^{\mathbf{e}_{S_\iota} - \mathbf{e}_T} \cdot \mathbf{A}[(k, \iota - \mathbf{e}_T), :] \cdot \mathbf{Z}^{\sum_{\ell \in S_\iota} d^{m-\ell}}. \quad (40)$$

By (39), we have that for every $\iota = (\iota^0, \iota^1)$ such that $\iota^1 \neq \mathbf{0}$

$$E[(k, \iota), J] = 0.$$

In other words, the non-zero rows (k, ι) must satisfy $\iota^1 = \mathbf{0}$. Thus, we have argued that

Lemma 10.6. *\mathbf{E} as defined in (40) has rank at most $nr^{m'}$.*

Proof. This follows from the fact that there are $n \cdot r^{m'}$ many indices (k, ι) with $\iota^1 = \mathbf{0}$. □

This in turn implies the following:

Lemma 10.7. *The recurrence in (40) has recurrence width of $\left(t = \frac{r^{m-m'+1}-1}{r-1}, nr^{m'}\right)$ and is $\left(D = \frac{d^{m-m'+1}-1}{d-1}, D\right)$ nice.*

10.2.1 Instantiation of parameters

We now consider few instantiation of parameters to get a feel for Lemma 10.7.

We start with the case of $n = 1$: note that in this case we have $r = d$ (since we want $N = n \cdot r^m = d^m$). In this case the choice of m' that makes all the parameters roughly equal is $m' = m/2$. In this case we get that (10.7) has recurrence width $(2d^{m/2}, d^{m/2})$ and is $(2d^{m/2}, 2d^{m/2})$ nice. However, this does not give anything algorithmically interesting since the input size for such a recurrence is already $\Omega(d^{m/2} \cdot d^{m/2} \cdot N + d^{m/2} \cdot N) = \Omega(N^2)$.

Recall that a recurrence with recurrence width (T, R) that is (D, D) nice has an input size of $O((TD + R)N)$. Thus, to decrease the input size of the recurrence in (40), we need to pick m' such that $2(m - m') = m'$. This implies that we pick $m' = 2m/3$ and thus we have a recurrence with recurrence width $(2d^{m/3}, 2d^{m/3})$ that is $(2d^{m/3}, 2d^{m/3})$ nice for an overall input size of $O(d^{5m/6}) = O(N^{5/3})$.

References

- [1] https://en.wikipedia.org/wiki/Approximation_theory#Chebyshev_approximation.
- [2] https://en.wikipedia.org/wiki/Jacobi_polynomials#Differential_equation.
- [3] https://en.wikipedia.org/wiki/Zernike_polynomials.
- [4] <http://wis.kuleuven.be/events/OPSFA/>.
- [5] <http://www.chebfun.org>.
- [6] NIST Handbook of Mathematical Functions. Cambridge University Press, New York, NY, 2010, ch. 24. Print companion to [16].
- [7] AHO, A. V., HOPCROFT, J. E., AND ULLMAN, J. D. *The Design and Analysis of Computer Algorithms*. Addison-Wesley, 1974.
- [8] APOSTOL, T. M. *Introduction to analytic number theory*. Springer Science & Business Media, 2013.
- [9] BELLA, T., EIDELMAN, Y., GOHBERG, I., AND OLSHEVSKY, V. Computations with quasiseparable polynomials and matrices. *Theoretical Computer Science* 409, 2 (2008), 158 – 179. Symbolic-Numerical Computations.
- [10] BINI, D., AND PAN, V. Y. *Polynomial and Matrix Computations (Vol. 1): Fundamental Algorithms*. Birkhauser Verlag, Basel, Switzerland, Switzerland, 1994.
- [11] BRENT, R. P., AND KUNG, H. T. Fast algorithms for manipulating formal power series. *J. ACM* 25, 4 (Oct. 1978), 581–595.
- [12] CANTOR, D. G., AND KALTOFEN, E. On fast multiplication of polynomials over arbitrary algebras. *Acta Informatica* 28, 7 (1991), 693–701.
- [13] CAO, Z., AND CAO, H. On fast division algorithm for polynomials using newton iteration. In *International Conference on Information Computing and Applications* (2012), Springer, pp. 175–180.
- [14] CHIHARA, T. *An Introduction to Orthogonal Polynomials*. Dover Books on Mathematics. Dover Publications, 2011.
- [15] COOLEY, J. W., AND TUKEY, J. W. An algorithm for the machine calculation of complex fourier series. *Math. Comput.* 19 (1965), 297–301.
- [16] NIST Digital Library of Mathematical Functions. <http://dlmf.nist.gov/>, Release 1.0.11 of 2016-06-08. Online companion to [6].

- [17] DRISCOLL, J. R., HEALY, JR., D. M., AND ROCKMORE, D. N. Fast discrete polynomial transforms with applications to data analysis for distance transitive graphs. *SIAM J. Comput.* 26, 4 (Aug. 1997), 1066–1099.
- [18] DUMMIT, D. S., AND FOOTE, R. M. *Abstract algebra*, vol. 3. Wiley Hoboken, 2004.
- [19] FIDUCCIA, C. M. An efficient formula for linear recurrences. *SIAM J. Comput.* 14, 1 (1985), 106–112.
- [20] GERASOULIS, A. A fast algorithm for the multiplication of generalized hilbert matrices with vectors. *Mathematics of Computation* 50, 181 (1988), 179–188.
- [21] GIORGI, P. On polynomial multiplication in chebyshev basis. *IEEE Trans. Comput.* 61, 6 (June 2012), 780–789.
- [22] GOHBERG, I., AND OLSHEVSKY, V. Complexity of multiplication with vectors for structured matrices. *Linear Algebra and its Applications* 202 (1994), 163 – 192.
- [23] GURUSWAMI, V., AND SUDAN, M. Improved decoding of Reed-Solomon and algebraic-geometry codes. *IEEE Transactions on Information Theory* 45, 6 (1999), 1757–1767.
- [24] HARTMAN, P. A lemma in the theory of structural stability of differential equations. *Proceedings of the American Mathematical Society* 11, 4 (1960), 610–620.
- [25] HARVEY, D. A multimodular algorithm for computing bernoulli numbers. *Math. Comput.* 79, 272 (2010), 2361–2370.
- [26] HIGHAM, N. J. *Functions of matrices: theory and computation*. Siam, 2008.
- [27] HORN, R. A., AND JOHNSON, C. R., Eds. *Matrix Analysis*. Cambridge University Press, New York, NY, USA, 1986.
- [28] KAILATH, T., KUNG, S.-Y., AND MORE, M. Displacement ranks of matrices and linear equations. *Journal of Mathematical Analysis and Applications* 68, 2 (1979), 395 – 407.
- [29] KAILATH, T., AND SAYED, A. H. Displacement structure: Theory and applications. *SIAM Review* 37, 3 (1995), 297–386.
- [30] KALTOFEN, E., AND SAUNDERS, B. D. On wiedemann’s method of solving sparse linear systems. In *International Symposium on Applied Algebra, Algebraic Algorithms, and Error-Correcting Codes* (1991), Springer Berlin Heidelberg, pp. 29–38.
- [31] KEDLAYA, K. S., AND UMANS, C. Fast polynomial factorization and modular composition. *SIAM J. Comput.* 40, 6 (2011), 1767–1802.
- [32] KOPPARTY, S. List-decoding multiplicity codes. *Theory of Computing* 11 (2015), 149–182.
- [33] KOPPARTY, S., MEIR, O., RON-ZEWI, N., AND SARAFA, S. High rate locally-correctable and locally-testable codes with sub-polynomial query complexity. *CoRR abs/1504.05653* (2015).
- [34] KOPPARTY, S., SARAFA, S., AND YEKHANIN, S. High-rate codes with sublinear-time decoding. *J. ACM* 61, 5 (2014), 28:1–28:20.
- [35] MILLER, K. S. On linear difference equations. *The American Mathematical Monthly* 75, 6 (1968), 630–632.
- [36] NÜSKEN, M., AND ZIEGLER, M. Fast multipoint evaluation of bivariate polynomials. In *Algorithms - ESA 2004, 12th Annual European Symposium, Bergen, Norway, September 14-17, 2004, Proceedings* (2004), pp. 544–555.
- [37] OLSHEVSKY, V., AND SHOKROLLAHI, M. A. A displacement approach to efficient decoding of algebraic-geometric codes. In *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing, May 1-4, 1999, Atlanta, Georgia, USA* (1999), pp. 235–244.

- [38] OLSHEVSKY, V., AND SHOKROLLAHI, M. A. Matrix-vector product for confluent cauchy-like matrices with application to confluent rational interpolation. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, May 21-23, 2000, Portland, OR, USA* (2000), pp. 573–581.
- [39] PAN, V. Y. *Structured Matrices and Polynomials: Unified Superfast Algorithms*. Springer-Verlag New York, Inc., New York, NY, USA, 2001.
- [40] PAN, V. Y., AND TSIGARIDAS, E. P. Nearly optimal computations with structured matrices. In *Symbolic-Numeric Computation 2014, SNC '14, Shanghai, China, July 28-31, 2014* (2014), pp. 21–30.
- [41] PUTZER, E. J. Avoiding the jordan canonical form in the discussion of linear systems with constant coefficients. *The American Mathematical Monthly* 73, 1 (1966), 2–7.
- [42] TREFETHEN, L. N., AND BAU III, D. *Numerical linear algebra*, vol. 50. Siam, 1997.
- [43] VANDEBRIL, R., BAREL, M. V., GOLUB, G., AND MASTRONARDI, N. A bibliography on semiseparable matrices*. *CALCOLO* 42, 3, 249–270.
- [44] YAP, C.-K. *Fundamental problems of algorithmic algebra*, vol. 49. Oxford University Press Oxford, 2000.
- [45] YU, K.-H., ZHANG, C., BERRY, G. J., ALTMAN, R. B., RÉ, C., RUBIN, D. L., AND SNYDER, M. Predicting non-small cell lung cancer prognosis by fully automated microscopic pathology image features. *Nature Communications* 7 (2016).
- [46] ZEILBERGER, D. A holonomic systems approach to special functions identities. *Journal of Computational and Applied Mathematics* 32, 3 (1990), 321 – 368.
- [47] ZHLOBICH, P., BELLA, T., EIDELMAN, Y., GOHBERG, I., AND OLSHEVSKY, V. *Classifications of Recurrence Relations via Subclasses of (H,m) -quasiseparable Matrices*, vol. 15 of *Lecture Notes in Electrical Engineering*. Springer-Verlag GmbH, 2011, p. 23.

A Related Work and Known Results

Superfast structured matrix vector multiplication has been a rich area of research. Toeplitz, Hankel, and Vandermonde matrices and their inverses all have classical superfast multiplication algorithms that correspond to operations on polynomials [7, 10]. A superfast algorithm was developed for Cauchy matrices (and slight generalizations) was developed by Gerasoulis in 1988 [20]. Multiplication with Cauchy matrices corresponds to operations on rational polynomials; Cauchy matrices naturally fit in with the other three types of matrices. The four classes of matrices are all generalized by the notion of displacement rank introduced by Kailath, Kung, and Morf in 1979 [28]. Kailath et al. used displacement rank to define Toeplitz-like matrices, generalizing Toeplitz matrices. In 1994, Gohberg and Olshevsky further used displacement rank to define Vandermonde-like, Hankel-like, and Cauchy-like matrices and developed superfast algorithms for vector multiplication [22]. These matrix classes were unified and generalized by Olshevsky and Shokrollahi in 2000 [38] with a class of matrices they named confluent Cauchy-like. Confluent Cauchy-like matrices are those with low displacement rank with respect to Jordan form matrices; we extend these results by investigating matrices with low displacement rank with respect to any triangular Δ -band matrices, which we define to be matrices whose non-zero elements all appear in Δ consecutive diagonals. Algorithms for these four classes of matrices have continued to be refined for precision; we refer to Pan and Tsigaridas in 2014 for an overview and for an algorithm with explicit bounds on precision [40].

Our work is spiritually closer to the study of orthogonal polynomial transforms, especially that of Driscoll, Healy, and Rockmore [17]. Orthogonal polynomials are widely used and well worth studying in their own right: for an introduction to the area, see the classic book of Chihara [14]. We present applications for some specific orthogonal polynomials. Chebyshev polynomials are used for numerical stability (see e.g. the ChebFun package [5]) as well as approximation theory (see e.g. Chebyshev approximation [1]). Jacobi polynomials form solutions of certain differential equations [2]. Zernike polynomials have applications in optics and physics [3]. In fact, our

investigation into structured matrix-vector multiplication problems started with some applied work on Zernike polynomials, and our results applied to fast Zernike transforms have been used in improved cancer imaging [45]. Driscoll et al. rely heavily on the three-term recurrence satisfied by orthogonal polynomials to devise a divide-and-conquer algorithm for computing matrix-vector multiplication. Our first main result is a direct generalization of the recurrence, and we rely heavily on the recurrence to formulate our own divide and conquer algorithm. As discussed earlier, we view the connection of these two strands of work as our strongest conceptual contribution.

A third significant strand of research is the study of semiseparable matrices. We refer to an extensive survey by Vandebril, Van Barel, Golub, and Mastronardi [43] for a detailed discussion of the body of work, but we provide a brief commentary here. The most straightforward class of semiseparable matrices are the generator representable semiseparable matrices, which are matrices whose upper triangular and lower triangular portions are both of (low) rank. Our results can straightforwardly recover generator semiseparable matrices whose triangular portions are of rank 1. However, semiseparable matrices are defined more generally with respect to the rank of matrix sub-blocks. The idea of utilizing the ranks of matrix sub-blocks has been generalized many times, and we refer to Bella, Eidelman, Gohberg, and Olshevsky's work on (H, m) -quasiseparable matrices [47] for a relatively recent exploration of various generalizations. This generalization actually has deep connections to polynomials that satisfy recurrences; if we define $p_i(X)$ as the characteristic polynomial of the upper left $i \times i$ submatrix, the p_i form a family that satisfy recurrence relations. Connecting our notion of recurrence - where the rows of our matrix form a recursive polynomial family - to that of quasiseparable matrices will be explored in future work. As far as we are aware, there is no fully general superfast matrix-vector multiplication algorithm for (H, m) -quasiseparable matrices. We note that this is an area of great active research; we point the reader to Bella et al.'s survey on computing with quasiseparable matrices [9] for an overview of the wide array of work.

A.1 Known Results

Multiplication of two degree N univariate polynomials can be performed in $\tilde{O}(N)$ operations: The classic FFT computes it in $O(N \log N)$ operations for certain fields [15], and generalizations of the Schonhage-Strassen algorithm compute it in $O(N \log N \log \log N)$ ring operations in general [12]. Computing polynomial divisors and remainders, i.e. $p(X) \pmod{q(X)}$ with $\deg(p(X)), \deg(q(X)) = O(N)$ can be done with the same number of operations [13].

The matrix-vector product by matrices \mathbf{A}, \mathbf{A}^T (and $\mathbf{A}^{-1}, \mathbf{A}^{-T}$ when they exist) for Toeplitz and Hankel matrices \mathbf{A} can be computed in $O(N \log N)$ operations [39]. When \mathbf{A} is a Vandermonde matrix, these products takes $O(N \log^2 N)$ operations [39]. The same holds for confluent Vandermonde matrices, defined as

Definition A.1. Given a set of points p_0, \dots, p_{N-1} , let $n_i = \sum_{j=0}^{i-1} \mathbb{1}(p_j = p_i)$ be the number of preceding points identical to p_i . Then the $N \times M$ (confluent) Vandermonde matrix⁹, denoted $\mathbf{V}_{p_0, \dots, p_{N-1}}$, is defined such that

$$\mathbf{V}[i, j] = \begin{cases} 0, & j < n_i \\ \frac{j!}{(j-n_i)!(n_i)!} p_i^{j-n_i}, & j \geq n_i \end{cases}$$

for $0 \leq i < N, 0 \leq j < M$.¹⁰

Furthermore, multiplication by these matrices encode polynomial evaluation. For any vector \mathbf{y} , $\mathbf{V}_{p_0, \dots, p_{N-1}} \mathbf{y}$ is equivalent to evaluating the polynomial $v(X) = \sum_{i=0}^{N-1} \mathbf{y}[i] X^i$ at $v^{(n_i)}(p_i)$ for each i .

The characteristic polynomial $c_{\mathbf{M}}(X)$ of a matrix \mathbf{M} is equal to the determinant of $X\mathbf{I} - \mathbf{M}$. When \mathbf{M} is triangular, it is equal to $\prod (X - \mathbf{M}[i, i])$. By the Cayley-Hamilton Theorem, every matrix satisfies its own characteristic equation, i.e. $c_{\mathbf{M}}(\mathbf{M}) = 0$.

⁹The term Vandermonde is usually applied when the p_i are all distinct, otherwise it is known as a confluent Vandermonde matrix. However, either way the matrix is uniquely specified by the ordered sequence p_i so we drop this distinction when the context is clear

¹⁰Some sources define the confluent Vandermonde matrix with the factor of $n_i!$ in the denominator [39], and others do not. It makes no real difference, but is slightly more convenient for us to use the former notation.

A λ -*Jordan block* is a square matrix of the form $\lambda \mathbf{I} + \mathbf{S}$, where \mathbf{S} is the shift matrix which is 1 on the superdiagonal and 0 elsewhere. A matrix in *Jordan normal form* is a direct sum of Jordan blocks. The minimal polynomial of a matrix is equal to the product $\prod_{\lambda} (X - \lambda)^{n_{\lambda}}$ where n_{λ} is the size of the largest Jordan block for λ . Consequently, if a matrix has equal minimal and characteristic polynomials, then it has only one Jordan block per eigenvalue.

B Examples of Recurrence Width in Matrices

For concreteness, we provide a few typical cases of matrices with low recurrence width and show how they fall under the above definitions.

Orthogonal Polynomials Consider the Chebyshev polynomials defined by $f_0(X) = 1, f_1(X) = X, f_{i+1}(X) = 2Xf_i(X) - f_{i-1}(X)$. The orthogonal polynomial transform matrix with respect to a set of points $\{z_0, \dots, z_{N-1}\}$ is

$$\mathbf{A} = \begin{bmatrix} f_0(z_0) & \cdots & f_0(z_{N-1}) \\ \vdots & \ddots & \vdots \\ f_{N-1}(z_0) & \cdots & f_{N-1}(z_{N-1}) \end{bmatrix}$$

The rows of this matrix, $\mathbf{f}_i = \mathbf{A}[i, :]^T$, satisfy $\mathbf{f}_0 = [1 \ \cdots \ 1]^T$, $\mathbf{f}_1 = [z_0 \ \cdots \ z_{N-1}]$ and $\mathbf{f}_{i+1} = 2\mathbf{R}\mathbf{f}_i - \mathbf{f}_{i-1}$. Thus this matrix has recurrence width 1 under Definition 2.4 with $g_{i,0}(X) = 2X$, $g_{i,1}(X) = -1$, $\mathbf{R} = \text{diag}(z_0, \dots, z_{N-1})$. More generally, for the orthogonal polynomial transform where the polynomials $f_i(X)$ satisfy the recurrence 1 and $\mathbf{A}[i, j] = f_i(z_j)$. In this case if define $\mathbf{f}_i = \mathbf{A}[i, :]^T$, then we can re-state (1) in our language as follows:

$$\mathbf{f}_{i+1} = a_i \mathbf{R} \mathbf{f}_i + b_i \mathbf{f}_i + c_i \mathbf{f}_{i-1}.$$

Low displacement rank matrices Toeplitz, Vandermonde and Cauchy matrices have low displacement rank, but their interpretation in our language is not that interesting because they define a generate recurrence (in particular, the recurrence for \mathbf{f}_i only has terms that depend on \mathbf{f}_i).

As another example we show the Pascal matrix, which is defined as $\mathbf{A}[i, j] = \mathbf{A}[i-1, j] + \mathbf{A}[i, j-1]$. Then one can show that

$$\begin{bmatrix} 0 & 0 & \cdots & 0 & 0 \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \end{bmatrix} \mathbf{A} - \mathbf{A} \begin{bmatrix} 1 & -1 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & -1 & 0 \\ 0 & 0 & \cdots & 1 & -1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & 1 \end{bmatrix} = \begin{bmatrix} -1 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & 0 \end{bmatrix}$$

so that it has displacement rank 1 with respect to Jordan matrices \mathbf{L}, \mathbf{R} , so it is a confluent Cauchy-like matrix [38]. Isolating the $i+1$ th row yields $\mathbf{f}_i^T - \mathbf{f}_{i+1}^T \mathbf{R} = \mathbf{0}^T$, so $\mathbf{f}_{i+1} = \mathbf{R}^{-T} \mathbf{f}_i$. In Appendix B, we show how this can be viewed as an instance of (11) under Definition 2.4.

Finally, we present a well-known matrix that is captured by our notion of recurrence but not by displacement rank. Consider the problem of evaluating a bivariate polynomial $f(X, Y) = \sum_{i=0}^{d-1} \sum_{j=0}^{d-1} f_{i,j} X^i Y^j$ on N points $(x_1, y_1), \dots, (x_N, y_N)$. Now consider the $N \times N$ matrix

$$\mathbf{A} = \begin{pmatrix} 1 & x_1 & \cdots & x_1^{d-1} & y_1 & y_1 x_1 & \cdots & y_1 x_1^{d-1} & y_1^2 & y_1^2 x_1 & \cdots & y_1^2 x_1^{d-1} & \cdots & y_1^{d-1} & y_1^{d-1} x_1 & \cdots & y_1^{d-1} x_1^{d-1} \\ 1 & x_2 & \cdots & x_2^{d-1} & y_2 & y_2 x_2 & \cdots & y_2 x_2^{d-1} & y_2^2 & y_2^2 x_2 & \cdots & y_2^2 x_2^{d-1} & \cdots & y_2^{d-1} & y_2^{d-1} x_2 & \cdots & y_2^{d-1} x_2^{d-1} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_N & \cdots & x_N^{d-1} & y_N & y_N x_N & \cdots & y_N x_N^{d-1} & y_N^2 & y_N^2 x_N & \cdots & y_N^2 x_N^{d-1} & \cdots & y_N^{d-1} & y_N^{d-1} x_N & \cdots & y_N^{d-1} x_N^{d-1} \end{pmatrix}.$$

Note that to solve the multipoint evaluation problem we just need to solve $\mathbf{A} \cdot \mathbf{f}$, where \mathbf{f} contains the coefficients of $f(X, Y)$. Let \mathbf{D}_X and \mathbf{D}_Y denote the diagonal matrices with $\mathbf{x} = (x_1, \dots, x_N)$ and $\mathbf{y} = (y_1, \dots, y_N)$ on their diagonals respectively. Finally, define $\mathbf{Z} = \mathbf{S}^T$.

We show in Section 10.1.1 that

$$\mathbf{D}_Y^{-1} \mathbf{D}_X^{-1} \mathbf{A} - \mathbf{D}_Y^{-1} \mathbf{A} \mathbf{Z} - \mathbf{D}_X^{-1} \mathbf{A} \mathbf{Z}^d - \mathbf{A} \mathbf{Z}^{d+1} = \mathbf{C}, \quad (41)$$

where

$$\mathbf{C} = \begin{pmatrix} \frac{1}{x_1 y_1} & 0 & \cdots & 0 & \frac{-x_1^{d-1}}{y_1} & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & \cdots & 0 & 0 & \cdots & 0 \\ \frac{1}{x_2 y_2} & 0 & \cdots & 0 & \frac{-x_2^{d-1}}{y_2} & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & \cdots & 0 & 0 & \cdots & 0 \\ & & & & \vdots & & & & & & & & \vdots & & & \\ \frac{1}{x_N y_N} & 0 & \cdots & 0 & \frac{-x_N^{d-1}}{y_N} & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & \cdots & 0 & 0 & \cdots & 0 \end{pmatrix}$$

has rank two.

We provide one more instantiation of the \otimes operator that arises in applications. In some of our applications (for example displacement rank in Section 9.1), a matrix $\mathbf{R} \in \mathbb{F}^{N \times N}$ is fixed and vectors \mathbf{f}_i are defined through a recurrence such as

$$D_{i+1}(\mathbf{R}) \mathbf{f}_{i+1} = \sum_{j=0}^t g_{i,j}(\mathbf{R}) \mathbf{f}_{i-j}$$

where all $D_{i+1}(\mathbf{R})$ are invertible or equivalently $\gcd(D_i, c_{\mathbf{R}}) = 1$. This equation is well-defined by $\mathbf{f}_{i+1} = \sum (D_{i+1}(\mathbf{R})^{-1} g_{i,j}(\mathbf{R})) \mathbf{f}_{i-j}$, which can be rewritten as

$$\mathbf{f}_{i+1} = \sum_{j=0}^t (g_{i,j}/D_{i+1}) \otimes \mathbf{f}_{i-j}$$

under the following definition.

Definition B.1. Let $\mathbf{R} \in \mathbb{F}^{N \times N}$ and let \mathcal{R} be the subring of $\mathbb{F}(X)$ consisting of fractions whose denominators are not a multiple of $c_{\mathbf{R}}(X)$. Given $a(X) = (b(X)/c(X)) \in \mathcal{R}$, the evaluation at \mathbf{R} is naturally defined as $a(\mathbf{R}) = b(\mathbf{R})c(\mathbf{R})^{-1} = c(\mathbf{R})^{-1}b(\mathbf{R})$. Now define $a(X) \otimes \mathbf{z} = a(\mathbf{R})\mathbf{z}$.

However, note the following: Let $a(X), b(X) \in \mathbb{F}[X]$ such that $(b, c_{\mathbf{R}}) = 1$, and let $c(X) = a(X)b^{-1}(X) \pmod{c_{\mathbf{R}}(X)}$. It is straightforward to show that $a(\mathbf{R})/b(\mathbf{R}) = c(\mathbf{R})$ because $c_{\mathbf{R}}(\mathbf{R}) = 0$. Therefore it is also true that

$$\mathbf{f}_{i+1} = \sum_{j=0}^t (g_{i,j}/D_{i+1}) \otimes \mathbf{f}_{i-j}$$

under Definition 2.4. Importantly, here $g_{i,j}/D_{i+1}$ is treated as an element of $\mathbb{F}[X]/(c_{\mathbf{R}}(X))$ (in particular, it can be represented as a polynomial) instead of $\mathbb{F}(X)$ (where it is represented as a fraction). Thus Definition B.1 is actually equivalent to Definition 2.4, and when it arises we will automatically assume that the recurrence is interpreted under 2.4.

This is how we are formally defining recurrences such as (7): As written it considers matrices in a polynomial modulus which does not make sense, but we use it to represent defining a polynomial recurrence with coefficients in $\mathbb{F}[X]/(c_{\mathbf{R}}(X))$ and then evaluating at \mathbf{R} .

C Details on the Work-horse Lemma

Proof of Lemma 2.6. Note that (11) can be written as

$$\mathbf{f}_{i+1} = [g_{i,0}(X) \quad \cdots \quad g_{i,t}(X)] \otimes [\mathbf{f}_i \quad \cdots \quad \mathbf{f}_{i-t}]^T$$

or

$$[\mathbf{f}_{i+1} \quad \cdots \quad \mathbf{f}_{i-t+1}]^T = \mathbf{T}_i \otimes [\mathbf{f}_i \quad \cdots \quad \mathbf{f}_{i-t}]^T$$

By composing this, we get

$$\begin{aligned}
[\mathbf{f}_{k+i} \quad \cdots \quad \mathbf{f}_{k+i-t}]^T &= \mathbf{T}_{k+i-1} \otimes (\cdots \otimes (\mathbf{T}_k \otimes [\mathbf{f}_{k+t} \quad \cdots \quad \mathbf{f}_k]^T)) \\
&= (\mathbf{T}_{k+i-1} \cdot \mathbf{T}_{k+i-2} \cdots \mathbf{T}_k) \otimes [\mathbf{f}_{k+t} \quad \cdots \quad \mathbf{f}_k]^T \\
&= \mathbf{T}_{[k:k+i]} \otimes [\mathbf{f}_{k+t} \quad \cdots \quad \mathbf{f}_k]^T
\end{aligned}$$

The first part of the lemma is equivalent to the top row of this equation, where we define $h_{i,j}^{(k)}$ to be $(1, j)$ -th entry of $\mathbf{T}_{[k:k+i]}$. The second part of the lemma is equivalent to the top row of $\mathbf{T}_{[k:k+i+1]} = \mathbf{T}_{k+i} \mathbf{T}_{[k:k+i]}$. \square

Proof of Lemma 2.7. First we prove this when $(d, \bar{d}) = (1, 0)$. Fix any arbitrary ℓ . We will prove the statement by induction on $r - \ell$. For the base case (i.e. when we are considering \mathbf{T}_ℓ), the bounds follow from the definition of \mathbf{T}_ℓ and our assumption on the sizes of $g_{\ell,j}(X)$ for $0 \leq j \leq t$. In fact when $r \leq \ell + t$ it can be shown inductively that the first $r - \ell$ rows satisfy these degree constraints and the rest looks like a shift matrix, i.e. $\mathbf{T}_{[\ell:r]}[i, j] = \delta_{r-\ell+j-i}$ for $i \geq r - \ell$. Now assume the result is true for $r - \ell = \Delta \geq t + 1$.

Now consider the case of $r = \ell + \Delta + 1$. In this case note that $\mathbf{T}_{[\ell:r]} = \mathbf{T}_{r-1} \cdot \mathbf{T}_{[\ell:r-1]}$. Now by the action of \mathbf{T}_{r-1} , the last t rows of $\mathbf{T}_{[\ell:r]}$ are the first t rows of $\mathbf{T}_{[\ell:r-1]}$ and the size claims for entries in those rows follows from the inductive hypothesis. Now note that for any $0 \leq j \leq t$, we have

$$\mathbf{T}_{[\ell:r]}[0, j] = \sum_{k=0}^t g_{r-1,k}(X) \cdot \mathbf{T}_{[\ell:r-1]}[k, j].$$

By the inductive hypothesis, we have that

$$\deg \mathbf{T}_{[\ell:r]}[0, j] \leq \max_{0 \leq k \leq t} \deg g_{r-1,k} + \deg \mathbf{T}_{[\ell:r-1]}[k, j] \leq \max_k ((k+1) + (r-1-\ell+j-k)) = r - \ell + j,$$

as desired.

Next, consider a $(d, 0)$ -nice recurrence and say its transition matrices are \mathbf{T}'_i and $\mathbf{T}'_{[\ell:r]}$. Note that for any k, i, j , we have $\deg(\mathbf{T}'_k[i, j]) \leq d \deg(\mathbf{T}_k[i, j])$ by the definition of niceness. Since degrees are additive under multiplication, this implies that for all ℓ, r, i, j , $\deg(\mathbf{T}'_{[\ell:r]}[i, j]) \leq d \deg(\mathbf{T}_{[\ell:r]}[i, j])$ as desired.

Finally, consider a (d, \bar{d}) -nice recurrence. By Definition 2.5, we can write its transition matrices as $D_i(X)^{-1} \mathbf{T}''_i$ where $\deg(D_i(X)) \leq \bar{d}$ and $\deg(\mathbf{T}''_k[i, j]) \leq \bar{d} + \deg(\mathbf{T}'_k[i, j])$. Thus the product of transition matrices equals $(D_\ell(X) \cdots D_{r-1}(X))^{-1} \mathbf{T}''_{[\ell:r]}$ and by additivity of degrees $\deg(\mathbf{T}''_{[\ell:r]}[i, j]) \leq \bar{d}(r - \ell) + \deg(\mathbf{T}'_{[\ell:r]}[i, j])$. \square

D Algorithms for Multipoint Evaluation of Multivariate Polynomials

Here we recollect known algorithms for multipoint evaluation of multivariate polynomials. Recall the multipoint evaluation problem:

Definition D.1. Given an m -variate polynomial $f(X_1, \dots, X_m)$ such that each variable has degree at most $d - 1$ and $N = d^m$ distinct points $\mathbf{x}(i) = (x(i)_1, \dots, x(i)_m)$ for $1 \leq i \leq N$, output the vector $(f(\mathbf{x}(i)))_{i=1}^N$.

We will use the following reduction from [31, 36]. Let $\alpha_1, \dots, \alpha_N$ be distinct points. For $i \in [m]$, define the polynomial $g_i(X)$ of degree at most $N - 1$ such that for every $j \in [N]$, we have

$$g_i(\alpha_j) = x(j)_i.$$

Then as shown in [31], the multipoint evaluation algorithm is equivalent to computing the following polynomial:

$$c(X) = f(g_1(X), \dots, g_m(X)) \mod h(X),$$

where $h(X) = \prod_{j=1}^N (X - \alpha_j)$. In particular, we have $c(\alpha_i) = f(\mathbf{x}(i))$ for every $i \in [N]$. Thus, we aim to solve the following problem:

Definition D.2 (Modular Composition). Given a polynomial $f(X_1, \dots, X_m)$ with individual degree at most $d-1$ and $m+1$ polynomials $g_1(X), \dots, g_m(X), h(X)$ all of degree at most $N-1$ (where $N \stackrel{\text{def}}{=} d^m$), compute the polynomial

$$f(g_1(X), \dots, g_m(X)) \mod h(X).$$

As was noted in [36], if for the multipoint evaluation problem all the $x(j)_1$ for $j \in [N]$ are distinct, then we can take $\alpha_j = x(j)_1$ and in this case $\deg(g_1) = 1$ since we can assume that $g_1(X) = X$. We will see that this allows for slight improvement in the runtime. Also in what follows, we will assume that an $n \times n$ and $n \times n^2$ matrix can be multiplied with $O(n^{\omega_2})$ operations.

D.1 Algorithm for the general case

Consider Algorithm 3 (which is a straightforward generalization of the algorithm for $m = 1$ from [11]).

We will use X^ι for any $\iota = (i_1, \dots, i_m)$ to denote the monomial $\prod_{\ell=1}^m X_\ell^{i_\ell}$. Let k be any integer that divides d and define

$$q = \frac{d}{k}.$$

With this notation, write down f as follows

$$f(X_1, \dots, X_m) = \sum_{J \in \mathbb{Z}_q^m} \left(\sum_{\iota \in \mathbb{Z}_k^m} f_{J, \iota} \cdot X^\iota \right) \cdot X^{J \cdot k}, \quad (42)$$

where $f_{J, \iota}$ are constants.

Algorithm 3 Algorithm for Modular Composition: general case

Input: $f(X_1, \dots, X_m)$ in the form of (42) and $g_1(X), \dots, g_m(X), h(X)$ of degree at most $N-1$ with $N = d^m$

Output:

$$f(g_1(X), \dots, g_m(X)) \mod h(X)$$

- 1: Let k be an integer that divides d ▷ We will use $k = \sqrt{d}$
 - 2: $q \leftarrow \frac{d}{k}$
 - 3: **For** every $\iota = (i_1, \dots, i_m) \in \mathbb{Z}_k^m$ **do**
 - 4: $g_\iota(X) \leftarrow \prod_{\ell=1}^m (g_\ell(X))^{i_\ell} \mod h(X)$.
 - 5: **For** every $J = (j_1, \dots, j_m) \in \mathbb{Z}_q^m$ **do**
 - 6: $g^J(X) \leftarrow \prod_{\ell=1}^m (g_\ell(X))^{j_\ell \cdot k} \mod h(X)$.
 - 7: **For** every $J \in \mathbb{Z}_q^m$ **do**
 - 8: $a_J(X) \leftarrow \sum_{\iota \in \mathbb{Z}_k^m} f_{J, \iota} \cdot g_\iota(X)$
 - 9: **Return** $\sum_{J \in \mathbb{Z}_q^m} a_J(X) \cdot g^J(X) \mod h(X)$
-

Algorithm 3 presents the algorithm to solving the modular composition problem. The correctness of the algorithm follows from definition. We now argue its runtime.

Note that for a fixed $\iota \in \mathbb{Z}_k^m$, the polynomial $g_\iota(X)$ can be computed in $\tilde{O}(mN)$ operations since it involves m exponentiations and $m-1$ product of polynomials of degree at most $N-1 \mod h(X)$. Thus, Step 3 overall takes $\tilde{O}(m \cdot k^m \cdot N)$ many operations. By a similar argument Step 5 takes $\tilde{O}(m \cdot q^m \cdot N)$ operations. Step 9 needs q^m polynomial multiplication ($\mod h(X)$) and $q^m - 1$ polynomial multiplication where all polynomial are of degree at most $N-1$ and hence, this step takes $\tilde{O}(q^m \cdot N)$ operations. So all these steps overall take $\tilde{O}(m \cdot \max(k, q)^m \cdot d^m)$ many operations.

So the only step we need to analyze is Step 7. Towards this end note that for any $j \in \mathbb{Z}_q^m$

$$\begin{aligned} a_j(X) &= \sum_{i \in \mathbb{Z}_k^m} f_{j,i} \cdot g_i(X) \\ &= \sum_{i \in \mathbb{Z}_k^m} f_{j,i} \cdot \sum_{\ell=0}^{N-1} g_i[\ell] \cdot X^\ell \\ &= \sum_{\ell=0}^{N-1} \left(\sum_{i \in \mathbb{Z}_k^m} f_{j,i} \cdot g_i[\ell] \right) X^\ell. \end{aligned}$$

Thus, if we think of the $q^m \times d^m$ matrix \mathbf{A} , where $\mathbf{A}[j, \cdot]$ has the coefficients of $a_j(X)$, then we have

$$\mathbf{A} = \mathbf{F} \times \mathbf{G},$$

where \mathbf{F} is an $q^m \times k^m$ matrix with $F_{j,i} = f_{j,i}$ and \mathbf{G} is an $k^m \times d^m$ matrix with $G_{i,\ell} = g_i[\ell]$. Let $\omega(r, s, t)$ be defined so that one can multiply an $n^r \times n^s$ with an $n^s \times n^t$ matrix with $n^{\omega(r,s,t)}$ operations. If we set $k = d^\epsilon$ for some $0 \leq \epsilon \leq 1$, we have that Algorithm 3 can be implemented with

$$\tilde{O}(m \cdot d^{m(\max(\epsilon, 1-\epsilon)+1)} + d^{m \cdot \omega(1-\epsilon, \epsilon, 1)})$$

many operations. It turns out that the expression above is optimized at $\epsilon = \frac{1}{2}$, which leads to an overall (assuming m is a constant) $\tilde{O}(d^{\omega_2 m/2})$ many operations. Thus, we have argued that

Theorem D.3. *The modular composition problem with parameters d and m can be solved with $\tilde{O}(d^{\omega_2 m/2})$ many operations.*

D.1.1 A ‘direct’ algorithm for the multipoint evaluation case

We now note that one can convert Algorithm 3 into a “direct” algorithm for the multipoint evaluation problem. Algorithm 4 has the details.

Algorithm 4 Algorithm for Multipoint Evaluation

Input: $f(X_1, \dots, X_m)$ in the form of (42) and evaluation points $\mathbf{a}(i)$ for $i \in [N]$

Output:

$$(f(\mathbf{a}(i)))_{i \in [N]}$$

- 1: **For** every $\iota = (i_1, \dots, i_{m/2}) \in \mathbb{Z}_d^{m/2}$ **do**
 - 2: $\mathbf{g}_\iota \leftarrow (\prod_{\ell=1}^{m/2} (a(k)_\ell)^{i_\ell})_{k \in [N]}$
 - 3: **For** every $J = (j_1, \dots, j_m) \in \mathbb{Z}_d^{m/2}$ **do**
 - 4: $\mathbf{g}^J \leftarrow (\prod_{\ell=m/2+1}^m (a(k)_\ell)^{j_\ell})_{k \in [N]}$
 - 5: **For** every $J \in \mathbb{Z}_d^{m/2}$ **do**
 - 6: $\mathbf{b}_J \leftarrow (\sum_{\iota \in \mathbb{Z}_d^{m/2}} f_{J,\iota} \cdot \mathbf{g}_\iota(k))_{k \in [N]}$
 - 7: **Return** $\sum_{J \in \mathbb{Z}_d^{m/2}} \langle \mathbf{b}_J, \mathbf{g}^J \rangle$
-

The correctness of the algorithm again follows from definition. It is easy to check that the computation of $\mathbf{g}_\iota, \mathbf{g}^J$ and the output vectors can be accomplished with $\tilde{O}(d^{3m/2})$ many operations. Finally, the computation of the \mathbf{b}_J can be done with $\tilde{O}(d^{\omega_2 m/2})$ many operations using fast rectangular matrix multiplication.

D.2 Algorithm for the distinct first coordinate case

We now consider the case when all the $x(j)_1$ for $j \in [N]$ are distinct: i.e. we assume that $g_1(X) = X$. In this case we re-write f as follows:

$$f(X_1, \dots, X_m) = \sum_{J \in \mathbb{Z}_q^{m-1}} \left(\sum_{\iota \in \mathbb{Z}_k^{m-1}} f_{J,\iota}(X_1) \cdot X_{-1}^\iota \right) \cdot X_{-1}^{J \cdot k}, \quad (43)$$

where X_{-1}^ι denotes the monomial on the variables X_2, \dots, X_m and each $f_{J,\iota}(X_1)$ is of degree at most $d-1$.

Algorithm 5 Algorithm for Modular Composition: distinct first coordinate case

Input: $f(X_1, \dots, X_m)$ in the form of (43) and $g_2(X), \dots, g_m(X), h(X)$ of degree at most $N-1$ with $N = d^m$

Output:

$$f(X, g_2(X), \dots, g_m(X)) \mod h(X)$$

- 1: Let k be an integer that divides d ▷ We will use $k = \sqrt{d}$
 - 2: $q \leftarrow \frac{d}{k}$
 - 3: **For** every $\iota = (i_2, \dots, i_m) \in \mathbb{Z}_k^{m-1}$ **do**
 - 4: $g_\iota(X) \leftarrow \prod_{\ell=2}^m (g_\ell(X))^{i_\ell} \mod h(X)$.
 - 5: **For** every $J = (j_2, \dots, j_m) \in \mathbb{Z}_q^{m-1}$ **do**
 - 6: $g^J(X) \leftarrow \prod_{\ell=2}^m (g_\ell(X))^{j_\ell \cdot k} \mod h(X)$.
 - 7: **For** every $J \in \mathbb{Z}_q^{m-1}$ **do**
 - 8: $a_J(X) \leftarrow \sum_{\iota \in \mathbb{Z}_k^{m-1}} f_{J,\iota}(X) \cdot g_\iota(X) \mod h(X)$
 - 9: **Return** $\sum_{J \in \mathbb{Z}_q^{m-1}} a_J(X) \cdot g^J(X) \mod h(X)$
-

Algorithm 5 shows how to update Algorithm 3 to handle this special case. Again the correctness of this algorithm follows from the definitions.

We next quickly outline how the analysis of Algorithm 5 differs from that of Algorithm 3. First, the same argument we used earlier can be used to show that Steps 3, 5 and 9 can be accomplished with $\tilde{O}(m \cdot \max(k, q)^{m-1} \cdot d^m)$ many operations.

As before the runtime is dominated by the number of operations needed for Step 7. Towards this end note that

$$\begin{aligned} a_J(X) &= \sum_{\iota \in \mathbb{Z}_k^{m-1}} f_{J,\iota}(X) \cdot g_\iota(X) \\ &= \sum_{\iota \in \mathbb{Z}_k^{m-1}} f_{J,\iota}(X) \cdot \sum_{\ell=0}^{d^{m-1}-1} g_\iota[\ell](X) \cdot (X^d)^\ell \\ &= \sum_{\ell=0}^{d^{m-1}-1} \left(\sum_{\iota \in \mathbb{Z}_k^{m-1}} f_{J,\iota}(X) \cdot g_\iota[\ell](X) \right) (X^d)^\ell. \end{aligned}$$

Note that in the above we have decomposed g_ι as a polynomial in powers of X^d (instead of X for Algorithm 3). In particular, this implies that all of $f_{J,\iota}(X)$ and $g_\iota[\ell](X)$ are polynomials of degree at most $d-1$. If we think of $a_J(X)$ as polynomials in X^d (with coefficients being polynomials of degree at most $2d-2$), we can represent the above as

$$\mathbf{A} = \mathbf{F} \times \mathbf{G},$$

where the $q^{m-1} \times k^{m-1}$ matrix \mathbf{F} is defined by $F_{J,\iota} = f_{J,\iota}(X)$ and the $k^{m-1} \times d^{m-1}$ matrix \mathbf{G} is defined by $g_{\iota,\ell} = g_\iota[\ell](X)$. Again if we set $k = n^\epsilon$, then the run time of Algorithm 5 is given by (we use the fact that each multiplication and addition in $\mathbf{F} \times \mathbf{G}$ can be implemented with $\tilde{O}(d)$ operations):

$$\tilde{O}(m \cdot d^{(m-1)\max(\epsilon, 1-\epsilon)+m} + d \cdot d^{(m-1) \cdot \omega(1-\epsilon, \epsilon, 1)})$$

many operations. It turns out that the expression above is optimized at $\epsilon = \frac{1}{2}$, which leads to an overall (assuming m is a constant) $\tilde{O}(d^{1+\omega_2(m-1)/2})$ many operations. Thus, we have argued that

Theorem D.4. *The modular composition problem with parameters d and m for the case of $g_1(X) = X$ can be solved with $\tilde{O}(d^{1+\omega_2(m-1)/2})$ many operations.*